

CO-CLUSTERING FOR DIRECTED GRAPHS; THE STOCHASTIC CO-BLOCKMODEL AND A SPECTRAL ALGORITHM

KARL ROHE¹ AND BIN YU²

¹UNIVERSITY OF WISCONSIN MADISON

²UNIVERSITY OF CALIFORNIA BERKELEY

ABSTRACT. Communities of highly connected actors form an essential feature in the structure of several empirical directed and undirected networks. However, compared to the amount of research on clustering for undirected graphs, there is relatively little understanding of clustering in directed networks.

This paper extends the spectral clustering algorithm to directed networks in a way that co-clusters or bi-clusters the rows and columns of a graph Laplacian. Co-clustering leverages the increased complexity of asymmetric relationships to gain new insight into the structure of the directed network. To understand this algorithm and to study its asymptotic properties in a canonical setting, we propose the Stochastic Co-Blockmodel to encode co-clustering structure. This is the first statistical model of co-clustering and it is derived using the concept of stochastic equivalence that motivated the original Stochastic Blockmodel. Although directed spectral clustering is not derived from the Stochastic Co-Blockmodel, we show that, asymptotically, the algorithm can estimate the blocks in a high dimensional asymptotic setting in which the number of blocks grows with the number of nodes. The algorithm, model, and asymptotic results can all be extended to bipartite graphs.

1. INTRODUCTION

A diverse set of data sources that are characterized by interacting units or actors can be easily represented as a network or graph. Social networks, representing people who communicate with each other, are one example. There are two basic types of networks, directed and undirected. The relationships in directed graphs are asymmetric. For example, in a communication network, one person calls the other person. In an undirected graph, all the relationships are symmetric. Although the network literature has typically addressed undirected networks, assuming symmetric relationships is often a simplification. Several types of relationships are asymmetric. For example, directed relationships compose

Thank you to David Gleich for thoughtful questions and helpful references. This research was primarily conducted while Karl Rohe was a graduate student and postdoc at UC Berkeley. He was partially supported by an NSF VIGRE Graduate Fellowship and ARO grant W911NF-11-1-0114. Bin Yu is partially supported by NSF grants SES-0835531 (CDI), DMS-1107000, 0939370 CCF, and ARO grant W911NF-11-1-0114.

communication networks, citation networks, web graphs, and internet networks. Even networks that are often represented as undirected networks (e.g. Facebook friendships and road networks) are simplifications of an underlying directed network. In Facebook, each friendship is proposed by one of the friends and received by the other friend. This induces an asymmetry. Facebook users also interact asymmetrically (e.g. “posting on walls,” “liking posts,” and “sharing posts”). In road networks, one is often interested in the flow of traffic; anyone who has a reverse commute can confirm that traffic flows asymmetrically. In biochemical cellular networks, a relationship represents the flow of information and/or energy in the cell. These are causal graphs, and causality without direction is merely correlation. In a wide range of applications, directed networks more accurately represent the collected data and possibly the corresponding data generating mechanism.

Just as it is common to remove edge direction and study the resulting undirected network, the extant literature on clustering algorithms has focused almost exclusively on undirected networks and presumed symmetric extensions for directed networks. One way to symmetrize a graph is to add an opposing edge to each observed edge. So, if the data contains an edge from a to b , then impute an edge from b to a . Satuluri and Parthasarathy [2011] propose 3 additional ways of symmetrizing directed networks for the purpose of clustering. However, there is additional information contained in the edge direction. This paper proposes *co-clustering* the nodes with a spectral method. Co-clustering leverages the increased complexity of directed relationships to allow for a richer set of possible studies on the structure of directed graphs.

Co-clustering (a.k.a. bi-clustering) was first proposed in Hartigan [1972] for Euclidean data arranged in a matrix $M \in \mathbb{R}^{n \times d}$. Where standard clustering techniques cluster the rows of M into k_r clusters, co-clustering simultaneously clusters the rows of M into k_r clusters and clusters the columns of M into k_c clusters. In the past decade, co-clustering has since become an important data analytic technique in biological applications (e.g. Madeira and Oliveira [2004], Tanay et al. [2004], Tanay et al. [2005], Madeira et al. [2010]), text processing (e.g. Dhillon [2001], Bisson and Hussain [2008]), and natural language processing (e.g. Freitag [2004], Rohwer and Freitag [2004]). Banerjee et al. [2004] suggests three advantages of co-clustering over traditional clustering. First, if the d columns are clustered into $k_c \ll d$ clusters, then it is easier to interpret why the row clusters are different. This is because the central point of each row cluster can now be represented in \mathbb{R}^{k_c} instead of the much larger \mathbb{R}^d . Second, clustering both the rows and the columns dramatically reduces the number of parameters, implicitly providing statistical regularization. To see this, notice that most clustering algorithms imply an approximate matrix decomposition $M \approx L\mu$, where each row of M is replaced by the central point of the cluster to which it belongs. In $L\mu$, $L \in \{0, 1\}^{n \times k_r}$, $L_{ij} = 1$ if row i is in cluster j , $\mu \in \mathbb{R}^{k_r \times d}$, and the j th row of μ contains the central point of cluster j . Co-clustering further decomposes $\mu = \tilde{\mu}R^T$, where $\tilde{\mu} \in \mathbb{R}^{k_r \times k_c}$, $R \in \{0, 1\}^{d \times k_c}$, and $R_{uv} = 1$ if column u (of M) is in column-cluster v . Thus, $M \approx L\tilde{\mu}R^T$. By decomposing μ into $\tilde{\mu}R^T$,

co-clustering provides statistical regularization by reducing the number of parameters. Finally, because there are fewer parameters to optimize, co-clustering algorithms have the potential to be faster than traditional clustering algorithms. For example, the computational complexity of a standard k-means step is ndk_r and the computational complexity of one iteration of an analogous co-clustering algorithm is $O(nk_ck_r + dk_ck_r)$ [Banerjee et al., 2004].

The previous applications of co-clustering have involved matrices where the rows and columns index different sets of objects. For example, in text processing, the rows correspond to documents, and the columns correspond to words. Element i, j of this matrix denotes how many times word j appears in document i . This paper applies co-clustering to an asymmetric square matrix where the rows and columns index the same set of nodes. As such, each node i is in two types of clusters, one corresponding to row i and the other corresponding to column i . The i th row of this matrix identifies the outgoing edges for node i . The i th column of this matrix identifies the incoming edges for node i . So, two rows are in the same co-cluster if they send edges to several of the same nodes; two columns are in the same co-cluster if they receive edges from several of the same nodes. In this way, co-clustering with directed graphs can lead to different types of interpretations and insights than previous applications of co-clustering (where the rows and columns index different sets).

This paper proposes and studies a spectral co-clustering algorithm called DI-SIM. The name DI-SIM has three meanings. First, because DI-SIM co-clusters the nodes, it uses two distinct (but related) similarity measures between nodes: “the number of common parents” and “the number of common offspring.” In this sense, DI-SIM means two similarities. Second, DI- also denotes that this algorithm is specifically for *directed* graphs. In fact, co-clustering on symmetric graphs produces two sets of redundant clusters. This is because row i is equal to column i . Additionally, DI-SIM, pronounced “dice ‘em,” dices data into clusters.

The rest of this paper is organized as follows: Section 2 defines DI-SIM and briefly explains how it differs from related algorithms. To study the performance of DI-SIM and conceptualize why it is reasonable to co-cluster the nodes in a directed graph, Section 3 discusses the notion of “stochastic equivalence” in directed graphs and proposes the Stochastic Co-Blockmodel to encode co-clustering structure. This model and the notion of “stochastic equivalence” allow insight into the graph asymmetries and the type of structure that co-clustering could estimate. Section 4 studies the asymptotic properties of DI-SIM on the Stochastic Co-Blockmodel. Theorem 4.1 shows that under certain conditions, DI-SIM can correctly co-cluster (i.e. estimate the co-blockmodel membership) of most nodes, even in the high dimensional setting where the number of blocks grows with the number of nodes. Section 5 presents two simulations investigating the two conditions in Theorem 4.1. Section 6 concludes the paper.

2. DI-SIM; A CO-CLUSTERING ALGORITHM FOR DIRECTED GRAPHS

The next subsection gives a brief overview of the classical spectral algorithm for undirected graphs. This motivates DI-SIM. A more detailed account of spectral clustering with undirected graphs can be found in von Luxburg [2007].

2.1. Spectral clustering for undirected graphs. Spectral clustering is a popular and computationally feasible algorithm for clustering the nodes of an undirected graph. It has a rich history in the algebraic connectivity of graphs. Since the seminal work of Donath and Hoffman [1973] and Fiedler [1973], the algorithm has been rediscovered and reapplied in numerous different fields. The algorithm is popular because it provides a computationally tractable approximation of the Cheeger constant, which is NP hard to exactly compute [Chung, 1997, Shi and Malik, 2000, von Luxburg, 2007]. Because of this, computer scientists have found many applications for variations of spectral clustering, such as load balancing and parallel computations [Van Driessche and Roose, 1995, Hendrickson and Leland, 1995], partitioning circuits for very-large-scale integration design [Hagen and Kahng, 1992] and sparse matrix partitioning [Pothén et al., 1990]. Spielman and Teng [2007] and von Luxburg et al. [2008] give detailed histories of the algorithm.

Shi and Malik [2000] popularized the use of spectral clustering for image segmentation. Since then, the algorithm has received significant attention in various machine learning applications. In the machine learning literature, the problem setup for spectral clustering is slightly different than the network approach in this paper. That literature presumes the data points lie in a metric space and a graph is constructed based on some measure of similarity between the points in this metric space. Several researchers have assumed that the data points lie on (or close to) a manifold and studied the asymptotic properties of the adjacency matrix and the graph Laplacian [Belkin, 2003, Bousquet et al., 2004, Coifman and Lafon, 2006, Giné and Koltchinskii, 2006, Hein, 2006, Hein et al., 2006, Belkin and Niyogi, 2008, Goldberg et al., 2008, Ting et al., 2011]. Others have studied the the adjacency matrix, the graph Laplacian, and spectral clustering under different conditions [Koltchinskii and Giné, 2000, von Luxburg et al., 2008, Shi et al., 2009, Rohe et al., 2011].

To define the spectral clustering algorithm, let $G = (V, E)$ denote a graph, where V is a vertex set and E is an edge set. The vertex set $V = \{v_1, \dots, v_n\}$ contains vertices or nodes. These are the actors in the systems discussed above. We will refer to vertex v_i as node i . This paper considers unweighted, directed edges. So, the edge set E contains a pair (i, j) if there is an edge, or relationship, from node i to node j : $i \rightarrow j$. The edge set can be represented by the adjacency matrix $A \in \{0, 1\}^{n \times n}$:

$$(2.1) \quad A_{ij} = \begin{cases} 1 & \text{if } (i, j) \text{ is in the edge set} \\ 0 & \text{otherwise.} \end{cases}$$

The graph is undirected if $(i, j) \in E \implies (j, i) \in E$. The adjacency matrix of such a graph is symmetric. The graph is directed if there exists a pair of nodes $\{i, j\}$ such that $(i, j) \in E$ and $(j, i) \notin E$. In this situation, the adjacency matrix is asymmetric.

The graph Laplacian is a function of the adjacency matrix. It is fundamental to spectral graph theory and the spectral clustering algorithm [Chung, 1997, von Luxburg, 2007]. For symmetric adjacency matrix A , define the symmetric graph Laplacian $L^{(s)}$ and diagonal matrix D both elements of $\mathbb{R}^{n \times n}$ in the following way,

$$(2.2) \quad \begin{aligned} D_{ii} &= \sum_{\ell=1}^n A_{i\ell} \\ L_{ij}^{(s)} &= \frac{A_{ij}}{\sqrt{D_{ii}D_{jj}}} = [D^{-1/2}AD^{-1/2}]_{ij}. \end{aligned}$$

Some readers may be more familiar defining $L^{(s)}$ as $I - D^{-1/2}AD^{-1/2}$. For spectral clustering, the difference is immaterial because both definitions have the same eigenvectors. Below is one version of the spectral clustering algorithm.

Spectral clustering according to Shi and Malik [2000]

Input: Symmetric adjacency matrix $A \in \{0, 1\}^{n \times n}$, number of clusters k .

- (1) Find the eigenvectors $X_1, \dots, X_k \in \mathbb{R}^n$ corresponding to the k largest eigenvalues of $L^{(s)}$. $L^{(s)}$ is symmetric, so choose these eigenvectors to be orthonormal. Form the matrix $X = [X_1, \dots, X_k] \in \mathbb{R}^{n \times k}$ by putting the eigenvectors into the columns.^a
- (2) Treating each of the n rows in X as a point in \mathbb{R}^k , run k -means with k clusters. This creates k non-overlapping sets A_1, \dots, A_k with union $1, \dots, n$.

Output: A_1, \dots, A_k . This means that node i is assigned to cluster g if the i th row of X is assigned to A_g in step 2.

^a[Shi and Malik, 2000] do not compute the eigenvector matrix X using $L^{(s)}$. Instead, they compute the largest generalized eigenvectors to the generalized eigenproblem $(D - A)x = \lambda Dx$. These formulations are mathematically equivalent.

This algorithm requires a symmetric adjacency matrix. As such, the graph must be undirected. DI-SIM, the co-clustering algorithm proposed in this paper is similar to the algorithm above. However, to accommodate an asymmetric matrix DI-SIM replaces the eigendecomposition with the singular value decomposition (SVD). SVD is defined as follows.

Definition 1. Singular value decomposition (SVD) factorizes a matrix $M \in \mathbb{R}^{n \times d}$ ($n \geq d$) into the product of orthonormal matrices $U \in \mathbb{R}^{n \times d}$, $V \in \mathbb{R}^{d \times d}$ and a diagonal matrix $\Sigma \in \mathbb{R}^{d \times d}$ with nonnegative entries.

$$M = U\Sigma V^T.$$

The columns of U contain the left singular vectors. The columns of V contain the right singular vectors. The diagonal of Σ contains the singular values. If the matrix M is square and symmetric, then SVD is equivalent to eigendecomposition and $U = V$. In this way, SVD is a generalization of the eigendecomposition.

2.2. The DI-SIM algorithm. Similarly to the undirected spectral clustering algorithm, DI-SIM utilizes a graph Laplacian. To define a graph Laplacian $L \in \mathbb{R}^{n \times n}$ for directed graphs, first define the $n \times n$ diagonal matrices P and O .

$$(2.3) \quad \begin{aligned} P_{ii} &= \sum_k A_{ki} = \sum_k \mathbf{1}\{k \rightarrow i\} \\ O_{ii} &= \sum_k A_{ik} = \sum_k \mathbf{1}\{i \rightarrow k\} \\ L_{ij} &= \frac{A_{ij}}{\sqrt{O_{ii}P_{jj}}} = \frac{\mathbf{1}\{i \rightarrow j\}}{\sqrt{O_{ii}P_{jj}}} = [O^{-1/2}AP^{-1/2}]_{ij} \end{aligned}$$

P_{ii} is the number of nodes that send an edge to node i , or the number of parents to node i . Similarly, O_{ii} is the number of nodes to whom i sends an edge, or the number of offspring to node i . The DI-SIM algorithm is defined as follows.

DI-SIM

Input: Adjacency matrix $A \in \{0, 1\}^{n \times n}$, number of row-clusters k_r , number of column-clusters k_c .

- (1) Compute the singular value decomposition of $L = U\Sigma V^T$. Discard the columns of U and V that correspond to the $n - k$ smallest singular values in Σ , where $k = \min\{k_r, k_c\}$. Call the resulting matrices $U^k \in \mathbb{R}^{n \times k}$ and $V^k \in \mathbb{R}^{n \times k}$.
- (2) Cluster the columns of L by treating each row of V^k as a point in \mathbb{R}^k . Cluster these points into k_c clusters with k -means. Because each row of V^k corresponds to a node in the graph, the resulting clusters are clusters of the nodes.
- (3) Cluster the rows of L by performing step two on the matrix U^k with k_r clusters.

Output: The clusters from step 2 and 3.

This algorithm is a generalization of the undirected spectral clustering algorithm. If A is the adjacency matrix of an undirected graph, then A is symmetric and $P = O$. This implies that $L = L^{(s)}$ is a symmetric matrix. Where the undirected spectral clustering algorithm uses the eigenvectors of $L^{(s)}$, the DI-SIM algorithm above uses the singular vectors of L . However, when L is symmetric, the eigenvectors of L are equivalent to the singular vectors of L . In this way, DI-SIM applied to an undirected graph is equivalent to spectral clustering.

SVD is a matrix decomposition technique that, on symmetric matrices, is equivalent to the eigendecomposition. This decomposition is an essential ingredient in DI-SIM. We are

not the first to apply SVD in a graph setting. The next subsection briefly highlights the previous algorithms that employ SVD to explore the structure of graphs.

2.2.1. Related SVD methods. DI-SIM provides novel insights into the structure of directed graphs by co-clustering the nodes with the left and right singular vectors. Several other researchers have used SVD to explore and understand different features of networks. For example, Dhillon [2001] used SVD to co-cluster bipartite graphs, Kleinberg [1999] used SVD in a web search algorithm that was a precursor to Google’s PageRank algorithm, Hoff [2009] used SVD to fit a random effects model to directed network data, and Sussman et al. [2011] used SVD to cluster the nodes of a directed graph.

Dhillon [2001] suggested an algorithm similar to DI-SIM that was to be applied to bipartite graphs in which the rows and columns of L correspond to different entities (e.g. documents and words). As such, L is rectangular. However, the definition of L remains the same, and the top singular vectors of this matrix play a similar role in both Dhillon [2001] and DI-SIM. There are four important differences between our algorithms. First, in DI-SIM, the rows and columns of L index the same set of nodes. So, each node is clustered into two types of clusters. In Dhillon [2001], the rows and columns index different sets of objects. As such, the interpretation of the clusters is different. Second, Dhillon clusters the rows and columns into the same number of clusters. DI-SIM allows for $k_c \neq k_r$. Third, Dhillon uses far fewer singular vectors. To find k clusters, Dhillon uses $\ell = \lceil \log_2 k \rceil$ singular vectors ($\lceil x \rceil$ is the smallest integer greater than x). While this makes the algorithm much faster, the current paper suggests that the singular vectors between ℓ and k can contain additional information. Finally, where DI-SIM runs k-means on U^k and V^k separately, Dhillon [2001] runs k-means on the rows of the matrix

$$\begin{pmatrix} O^{-1/2}U^\ell \\ P^{-1/2}V^\ell \end{pmatrix},$$

where U^ℓ and V^ℓ are the left and right singular vectors corresponding to the top ℓ singular values.

Kleinberg [1999] proposed the concept of “hubs and authorities” for hyperlink-induced topic search (HITS). This algorithm that was a precursor to Google’s PageRank algorithm [Page et al., 1999]. To find authoritative sources of information, HITS “propose[d] an algorithmic formulation of authority, based on the relationship between a set of relevant authoritative pages and the set of ‘hub pages’ that join them together in the link structure” Kleinberg [1999]. To perform a web search, the HITS algorithm first initializes with a set of roughly 100 websites that could have been found with the standard text analysis algorithms in 1999. Second, it constructs the subgraph composed of all sites that link to or are linked from those 100 pages. Finally, it takes the top left singular vector and the top right singular vector of the resulting adjacency matrix. These left and right singular vectors respectively assign each website a hub-ness and authoritativeness score. HITS uses only

the top singular vectors for web search. DI-SIM uses the top several singular vectors for clustering.

Hoff [2009] applied SVD to directed graphs to estimate “sender-specific and receiver-specific latent nodal attributes.” He used the following random effects model,

$$P(A|Z) = \prod_{ij} \frac{\exp(\beta^T x_{ij} + Z_{ij})}{1 + \exp(\beta^T x_{ij} + Z_{ij})}$$

where x_{ij} contains observed covariates and Z_{ij} is an “unobserved latent effect.” The matrix Z is modeled as a low rank matrix using the SVD of residuals from the model fit without Z . In Hoff’s words, “this approach allows for the graphical description of a social network via the latent factors of the nodes, and provides a framework for the prediction of missing links in network data.” The left and right singular vectors estimate the “sender-specific” and the “receiver-specific traits.”

In a separate line of research, Sussman et al. [2011] studied how the singular vectors of a random adjacency matrix (generated from the standard Stochastic Blockmodel) can estimate the block memberships. Sussman et al. cluster the nodes of the network, they do not *co*-cluster the nodes. Further, their Stochastic Blockmodel does not encode the concept of co-clustering. They use the left and right singular vectors in a way that effectively ignores edge direction.

SVD has been used in other statistical methodology as well, such as correspondence analysis (CA) and canonical correlation analysis. In fact, DI-SIM normalizes the rows and columns in an identical fashion to CA. CA is a multivariate methodology with a long and storied existence. CA has similarities to principal components analysis, but it is applicable to categorical data. The methodology was first published in Hirschfeld [1935] and, like spectral clustering, it has been rediscovered and reapplied several times over [Guttman, 1959]. This method has become popular in ecology and has been refined in several different ways [Hill, 1979, Ter Braak, 1986]. It is particularly popular among French ecologists [Holmes, 2006]. While there are strong algorithmic connections to CA, the language and techniques of CA have not yet been applied to network data. This is a potentially fruitful area for further research.

2.3. Interpreting the singular vectors. This subsection examines how the singular vectors of L use the similarity measures “number of common parents” and “number of common offspring.” Recall that SVD expresses a matrix $M \in \mathbb{R}^{n \times d}$, as the product of three matrices, $M = U\Sigma V^T$. Lemma 2.1 shows how to compute the matrices U , V , and Σ , giving insight into the similarity measures used by DI-SIM. The lemma follows from Lemma 7.3.1 in Horn and Johnson [2005].

Lemma 2.1. *With SVD, $M = U\Sigma V^T$ for $M \in \mathbb{R}^{n \times d}$. The matrices U and V contain the eigenvectors to the symmetric, positive semi-definite matrices MM^T and $M^T M$ respectively. Both MM^T and $M^T M$ have the same eigenvalues and these values are contained in the diagonal of Σ^2 . If M is symmetric, U contains the eigenvectors of M and $U = V$.*

This implies that DI-SIM uses the eigenvectors of two matrices, $L^T L$ and LL^T [Kleinberg, 1999]. To understand these matrices, first look at $A^T A$ and AA^T . When A is the adjacency matrix of a directed graph, $A^T A$ and AA^T are two symmetric similarity matrices that correspond to “the number of common parents” and “the number of common offspring.”

$$(A^T A)_{ab} = \sum_x \mathbf{1}\{x \rightarrow a \text{ and } x \rightarrow b\} : \text{The number of common “parents.”}$$

$$(AA^T)_{ab} = \sum_x \mathbf{1}\{a \rightarrow x \text{ and } b \rightarrow x\} : \text{The number of common “offspring”}$$

These two similarity matrices are symmetric and easily interpretable. LL^T and $L^T L$ perform a similar task while down-weighting the contribution of high degree nodes.

$$(L^T L)_{ab} = \sum_x L_{xa} L_{xb} = \frac{1}{\sqrt{P_{aa} P_{bb}}} \sum_x \frac{\mathbf{1}\{x \rightarrow a \text{ and } x \rightarrow b\}}{O_{xx}}$$

$$(LL^T)_{ab} = \sum_x L_{ax} L_{bx} = \frac{1}{\sqrt{O_{aa} O_{bb}}} \sum_x \frac{\mathbf{1}\{a \rightarrow x \text{ and } b \rightarrow x\}}{P_{xx}}$$

The next sections explicate the difference between these two similarity measures by exploring the relationship between DI-SIM and “stochastic equivalence,” a concept discussed in Section 3. Using the concept of “stochastic equivalence,” Section 3 introduces a new Stochastic Co-Blockmodel. This new model will serve as a test bed to examine the two types of clusters produced by DI-SIM. Section 4, shows that DI-SIM can asymptotically estimate the block structure in the Stochastic Co-Blockmodel.

3. STOCHASTIC CO-BLOCKMODEL

In a directed network, the “number of common parents” between two actors a and b is the number of actors that send a directed relationship to both a and b . The “number of common offspring” between a and b is the number of actors that receive a directed relationship from both a and b . To study these similarity measures, this paper relates the observable measurements to two *model based* notions of similarity.

3.1. Stochastic equivalence, a model based similarity. The Stochastic Blockmodel [Holland et al., 1983] in social network analysis is motivated by a model based notion of similarity called stochastic equivalence. Holland et al. [1983] defines stochastic equivalence

thusly,¹ “We say two nodes a and b are stochastically equivalent if and only if the probability of any event about A is unchanged by interchanging nodes a and b .” This means that elements of A indexed by a are exchangeable with elements of A indexed by b . In an undirected network, two actors are stochastically equivalent if only if they connect to any third actor with equal probabilities:

$$P(a \leftrightarrow x) = P(b \leftrightarrow x) \forall x,$$

where $a \leftrightarrow x$ denotes the event that a and x are connected. In a directed network, the definition of Holland et al. [1983] implies that two nodes a and b are stochastically equivalent if and only if both of the following hold:

$$(3.1) \quad P(a \rightarrow x) = P(b \rightarrow x) \forall x \text{ and}$$

$$(3.2) \quad P(x \rightarrow a) = P(x \rightarrow b) \forall x$$

where $a \rightarrow x$ denotes the event that a sends an edge to x .

The Stochastic Blockmodel builds off the notion of stochastic equivalence. In the Stochastic Blockmodel, each node belongs to a cluster, or block, and two nodes in the same block are stochastically equivalent. Conversely, if two nodes are stochastically equivalent, then they are in the same block.

To align with the concept of co-clustering, this paper suggests a relaxation of the notion of stochastic equivalence for directed graphs. Instead of one type of stochastic equivalence which implies both Equations (3.1) and (3.2), one can relax this notion and allow two different types of stochastic equivalence. Two nodes a and b are stochastically equivalent senders if and only if Equation (3.1) holds. Two nodes a and b are stochastically equivalent receivers if and only if Equation (3.2) holds. These two concepts correspond to a model based notion of co-clusters and they are simultaneously represented in the new Stochastic Co-Blockmodel.

3.2. A statistical model of co-clustering in directed graphs. The Stochastic Blockmodel provides a model for a random network with k well defined blocks, or communities [Holland et al., 1983]. Each node in the Stochastic Blockmodel is a member of exactly one block, and any two nodes within the same block connect to any third node with an equal probability. That is, nodes in the same block are stochastically equivalent. Further, edges are statistically independent. The following is a definition of the classical Stochastic Blockmodel.

Definition 2. Define two nonrandom matrices, $Z \in \{0, 1\}^{n \times k}$ and $B \in [0, 1]^{k \times k}$. Each row of Z has exactly one 1, each column has at least one 1, and B is symmetric. Under the **Stochastic Blockmodel**, the adjacency matrix $A \in \{0, 1\}^{n \times n}$ is symmetric and random

¹Our notation has been substituted for the notation in Holland et al. [1983]

such that $E(A) = ZBZ^T$. Additionally, each edge is independent, so the probability distribution factors

$$P(A) = \prod_{i,j} P(A_{ij}).$$

If the graph is undirected, then the product is only taken over the set $i \leq j$.

In this definition, $Z_{ia} = 1$ if the i th node is a member of the a th block and B_{ab} is the probability of a connection from a node in the a th block to a node in the b th block. Define z_i to be the i th row of Z . Notice that if $z_i = z_j$, then the i th row of $E(A)$ is equal to the j th row of $E(A)$ and the i th column of $E(A)$ is equal to the j th column of $E(A)$: $z_i B Z^T = z_j B Z^T$ and $Z B z_i = Z B z_j$. So, two nodes in the same block are stochastically equivalent. By assuming an identifiability type assumption on B , the converse is also true. For example, it is sufficient for B to be full rank.

The Stochastic Co-Blockmodel is an extension of the Stochastic Blockmodel.

Definition 3. Define three nonrandom matrices, $Y \in \{0, 1\}^{n \times k_y}$, $Z \in \{0, 1\}^{n \times k_z}$ and $B \in [0, 1]^{k_y \times k_z}$. Each row of Y and each row of Z has exactly one 1 and each column has at least one 1. Under the **Stochastic Co-Blockmodel**, the adjacency matrix $A \in \{0, 1\}^{n \times n}$ is random such that $E(A) = YBZ^T$. Further, each edge is independent, so the probability distribution factors

$$P(A) = \prod_{i,j} P(A_{ij}).$$

In the Stochastic Blockmodel $E(A) = ZBZ^T$. In the Stochastic Co-Blockmodel, $E(A) = YBZ^T$. In this definition, Y and Z record two types of block membership which correspond to the two types of stochastic equivalence (Equations (3.1) and (3.2)).

Proposition 3.1. Under the Stochastic Co-Blockmodel, let y_i be the i th row of Y . If $y_i = y_j$, then nodes i and j are stochastically equivalent senders, Equation (3.1). Similarly, if $z_i = z_j$, then nodes i and j are stochastically equivalent receivers, Equation (3.2).

Proof. If $y_i = y_j$, then the i and j th rows of $E(A)$ are equal. Thus, nodes i and j send edges to any third node x with equal probability. This implies that nodes i and j are stochastically equivalent senders, Equation (3.1). Similarly, if $z_i = z_j$, then the i and j th columns of $E(A)$ are equal. Thus, nodes i and j receive edges from any third node x with equal probability. This implies that nodes i and j are stochastically equivalent receivers, Equation (3.2). \square

Proposition 3.1 shows that the Stochastic Co-Blockmodel encodes the two types of stochastic equivalence in the matrices Y and Z . These two matrices encode the two types of clusters that co-clustering estimates.

While the Stochastic Co-Blockmodel provides a model based notion of co-clustering, it can be thought of as a reparameterization of the Stochastic Blockmodel. If one starts with

a Stochastic Co-Blockmodel with k_y and k_z blocks, then one can create a traditional Stochastic Blockmodel with (at most) $k_y k_z$ blocks. In the traditional Stochastic Blockmodel, two nodes are in the same block if and only if they are stochastically equivalent parents and they are stochastically equivalent offspring. So, by adding a block for every combination of a Y -block and a Z -block, the Stochastic Blockmodel can parameterize the same model. Parameterizing a Stochastic Co-Blockmodel as a Stochastic Blockmodel creates a problem in estimating the matrix of probabilities B . Under the classical model, estimating this matrix requires estimating up to $(k_y k_z)^2$ probabilities. In the Co-Blockmodel, estimating B requires estimating $k_y k_z$ probabilities. In this sense, every Stochastic Co-Blockmodel is a Stochastic Blockmodel with fewer parameters and a unique interpretation.

A previous version of the Stochastic Blockmodel for directed networks, proposed in Wang and Wong [1987], allows for random variable pairs A_{ij} and A_{ji} to be statistically dependent; this is a fundamental difference between the Stochastic Co-Blockmodel and the previously defined Stochastic Blockmodels for directed graphs. Definition 3 does not allow for any stochastic dependence. In this sense, Wang and Wong [1987] is more general than Definition 3. That said, the goal of the Stochastic Co-Blockmodel is to encode co-clustering structure. Statistical independence is assumed only for mathematical ease.

This paper uses DI-SIM to estimate the partitions in Y and Z . We study this spectral algorithm and avoid the maximum likelihood estimator (MLE) for two reasons. First, DI-SIM is computationally feasible while the MLE is computationally intractable to compute. Second, various versions of spectral methods have proved useful in a wide range of applications and appear to give more reasonable answers than techniques which rely on discrete optimization [Leskovec et al., 2008].

The next section gives conditions under which DI-SIM can asymptotically estimate the block partitions in the matrices Y and Z . This implies that the two notions of stochastic equivalence relate to the two sets of singular vectors of L .

4. ASYMPTOTIC PERFORMANCE OF DI-SIM UNDER THE STOCHASTIC CO-BLOCKMODEL

Theorem 4.1 bounds the number of nodes that DI-SIM “misclusters” asymptotically. This demonstrates that the co-clusters from DI-SIM estimate the two types of block membership, one in matrix Y and the other in matrix Z , corresponding to the two types of stochastic equivalence.

In a diverse set of large empirical networks, the optimal clusters, as judged by a wide variety of graph cut objective functions, are not very large [Leskovec et al., 2008]. To account for this, the results below limit the growth of community sizes by allowing the number of communities to grow with the number of nodes. Previously, Rohe et al. [2011] and Choi et al. [2012] have studied this high dimensional setting for the undirected Stochastic Blockmodel.

Rigorous discussions of clustering require careful attention to identifiability. In the Stochastic Co-Blockmodel, the *order* of the columns of Y and Z are unidentifiable. This leads to difficulty in defining “misclustered.” Rohe et al. [2011] gives a reasonable and tractable definition of misclustered that this paper extends to co-clustering. Sections C.1 and C.2 (in the Appendix) motivate and present the retrofitted definition.

4.1. Main result. Define $\mathcal{A} = E(A)$ as the population version of the adjacency matrix A . Under the Stochastic Co-Blockmodel,

$$\mathcal{A} = YBZ^T.$$

Define population versions of O , P , and L all in $\mathbb{R}^{n \times n}$ as

$$(4.1) \quad \begin{aligned} \mathcal{O}_{jj} &= \sum_k \mathcal{A}_{kj} \\ \mathcal{P}_{ii} &= \sum_k \mathcal{A}_{ik} \\ \mathcal{L} &= \mathcal{P}^{-1/2} \mathcal{A} \mathcal{O}^{-1/2} \end{aligned}$$

where \mathcal{O} and \mathcal{P} are diagonal matrices. The following definitions provide for an alternative definition of \mathcal{L} . Define

$$\begin{aligned} D_y &= \text{diag}(BZ^T \mathbf{1}_n) \in \mathbb{R}^{k_y \times k_y} \\ D_z &= \text{diag}(\mathbf{1}_n^T YB) \in \mathbb{R}^{k_z \times k_z} \\ B^{\mathcal{L}} &= D_y^{-1/2} B D_z^{-1/2} \in [0, 1]^{k_y \times k_z}, \end{aligned}$$

where $\text{diag}(x)$ for $x \in \mathbb{R}^d$ is a diagonal matrix in $\mathbb{R}^{d \times d}$ with $\text{diag}(x)_{jj} = x_j$. The following is an alternative expression for \mathcal{L} ,

$$\mathcal{L} = YB^{\mathcal{L}}Z^T.$$

Define P_{max}^y to be the population of the largest block in Y .

$$(4.2) \quad P_{max}^y = \max_{j=1, \dots, k_y} (Y^T Y)_{jj}$$

Define $B_{\cdot j}^{\mathcal{L}}$ as the j th column of $B^{\mathcal{L}}$, and define

$$(4.3) \quad \gamma_z = \sqrt{P_{min}^y} \min_{i \neq j} \|B_{\cdot i}^{\mathcal{L}} - B_{\cdot j}^{\mathcal{L}}\|_2,$$

where P_{min}^y is the population of the smallest block in Y . The next theorem bounds the size of the set of Y -misclustered nodes $|\mathcal{M}_y|$ and the size of the set of Z -misclustered nodes $|\mathcal{M}_z|$. The definitions of $|\mathcal{M}_y|$ and $|\mathcal{M}_z|$ are presented in the Appendix, in Sections C.1 and C.2.

Theorem 4.1. *Suppose $A \in \mathbb{R}^{n \times n}$ is an adjacency matrix sampled from the Stochastic Co-Blockmodel with k_y left blocks and k_x right blocks. Assume $k_y \leq k_z$. Define \mathcal{L} as in*

(4.1). Define $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{k_y} > 0$ as the k_y nonzero singular values of \mathcal{L} . Define \mathcal{M}_y and \mathcal{M}_z as the sets of Y - and Z -misclustered nodes, as in (C.11) and (C.15). Define

$$\tau_n = \min_{i=1,\dots,n} (\min\{\mathcal{O}_{ii}, \mathcal{P}_{ii}\}) / n.$$

Define P_{max}^y as in (4.2) and γ_z as in (4.3). Assume there exists N such that for all $n > N$, $\tau_n^2 > 2/\log n$. If $n^{-1/2}(\log n)^2 = O(\sigma_{k_y})$, then

$$(4.4) \quad |\mathcal{M}_y| = o\left(\frac{(\log n)^2 P_{max}^y}{n\sigma_{k_y}^4 \tau_n^4}\right)$$

$$(4.5) \quad |\mathcal{M}_z| = o\left(\frac{(\log n)^2}{n\sigma_{k_y}^4 \tau_n^4 \gamma_z^2}\right)$$

A proof of Theorem 4.1 is contained in Appendix C.

The bound on $|\mathcal{M}_z|$ includes γ_z^{-2} . The next proposition makes an additional assumption to simplify this quantity.

Proposition 4.1. *Under the Stochastic Co-Blockmodel, if all expected in- and out-degrees are equal,*

$$\mathcal{O}_{ii} = \mathcal{P}_{jj} = n\tau_n \text{ for all } i \text{ and } j,$$

then

$$\gamma_z = \frac{\sqrt{P_{min}^y}}{n\tau_n} \min_{i \neq j} \|B_{\cdot i} - B_{\cdot j}\|_2,$$

where $B_{\cdot i}$ is the i th column of B .

Under the assumptions of Theorem 4.1 and Proposition 4.1,

$$(4.6) \quad |\mathcal{M}_z| = o\left(\frac{n(\log n)^2}{\sigma_{k_y}^4 \tau_n^2 P_{min}^y \min_{i \neq j} \|B_{\cdot i} - B_{\cdot j}\|_2^2}\right)$$

One might be concerned with the n in the numerator of the right hand side of 4.6. The end of this section discusses the proportion of misclustered nodes, \mathcal{M}_y/n and \mathcal{M}_z/n . After dividing by n , the P_{min}^y in the denominator can drive the bound to zero.

The bound for \mathcal{M}_z appears to exceed the bound for \mathcal{M}_y . In fact, if $k_y = k_z$, then \mathcal{M}_z can be bounded as in (4.4) with P_{max}^y replaced with an analogous quantity for Z .² There is only a problem when $k_y < k_z$. $\text{Rank}(\mathcal{L})$ is at most k_y . So, the singular value decomposition represents the data in k_y dimensions and the k-means steps for both the left and the right clusters are done in k_y dimensions. In estimating Y , there is one dimension in the singular vector representation for each of the k_y blocks. At the same time, the singular

²If $k_y = k_z$, then the definition of \mathcal{M}_z , in the appendix, should be changed to a definition analogous to \mathcal{M}_y . The current definition of \mathcal{M}_z in Equation C.15 is formulated to allow $k_z > \text{rank}(\mathcal{L})$.

value representation shoehorns the k_z blocks in Z into less than k_z dimensions. So, there is less space to separate each of the k_z clusters, obscuring the estimation of Z .

The two main assumptions of Theorem 4.1 are

- (1) $n^{-1/2}(\log n)^2 = O(\sigma_{k_n})$
- (2) eventually, $\tau_n^2 \log n > 2$.

The first assumption, a spectral gap condition, requires that the smallest nonzero singular value of \mathcal{L} is large enough. The second assumption ensures that the expected degree of each node grows sufficiently fast. If τ_n is constant, then the expected degree of each node grows linearly. The assumption $\tau_n^2 > 2/\log n$ is almost as restrictive. These two conditions imply the sufficient conditions needed to show that the top k_y singular vectors of L converge to the top k_y singular vectors of \mathcal{L} . That theorem is contained in the Appendix.

To understand the bound in Theorem 4.1, define the following toy model.

Definition 4. *The four parameter Stochastic Co-Blockmodel is a Stochastic Co-Blockmodel parameterized by $k \in \mathbb{N}$, $s \in \mathbb{N}$, $r \in (0, 1)$, and $p \in (0, 1)$ such that $p + r \leq 1$. The matrices $Y, Z \in \{0, 1\}^{n \times k}$ each contain s ones in each column and $B = pI_k + r\mathbf{1}_k\mathbf{1}_k^T$.*

In the four parameter Stochastic Co-Blockmodel there are k left- and right-blocks each with s nodes and the node partitions in Y and Z are not necessarily related. If $y_i = z_j$, then $P(i \rightarrow j) = p + r$. Otherwise, $P(i \rightarrow j) = r$.

Corollary 4.1. *Assume the four parameter Stochastic Co-Blockmodel, with r and p fixed and k growing with $n = ks$. Then,*

$$\sigma_k = \frac{1}{k(r/p) + 1},$$

where σ_k is the k th largest singular value of \mathcal{L} . If $k = O(n^{1/4}/\log n)$, then

$$|\mathcal{M}_y| + |\mathcal{M}_z| = o(k^3 \log^2 n).$$

Further, the proportion of nodes that are misclustered converges to zero,

$$(|\mathcal{M}_y| + |\mathcal{M}_z|)/n = o(n^{-1/4}).$$

The proof of Corollary 4.1 is contained in Appendix C.³

³Corollary 4.1 adopts a definition of \mathcal{M}_z that is analogous to \mathcal{M}_y because in the four parameter model $k_y = k_z$.

5. SIMULATIONS

The three simulations in this section address three different aspects of the theoretical results in Theorem 4.1. The first two simulations investigate τ and σ_k respectively. These simulations investigate DI-SIM's non-asymptotic sensitivity to these quantities under the four parameter Stochastic Co-Blockmodel. The third simulation compares DI-SIM's performance to its theoretical guarantees as k_z grows and k_y remains fixed.

k -means is the final step of DI-SIM. k -means often gets trapped in local optima. To avoid this problem, this section initializes the k -means centers to the rotated population singular vectors $Z\mu O$. In preparing these simulations, this initialization was particularly helpful in the more challenging simulations, where greater than 50 of the nodes are misclustered.

In the following simulations, any simulated graph that had a node with no in-edges or no out-edges was discarded. The entire graph was re-simulated until every in- and out-degree was nonzero. In this sense, every simulated graph is conditioned to have strictly positive in- and out-degrees.

5.1. Simulations 1 and 2. The first two simulations come from the four parameter model with 30 blocks ($k = 30 = k_y = k_z$) and 30 nodes in each block ($s = 30$). In the first set of simulations, the probabilities p and r vary in such a way that τ decreases and σ_k remains fixed. In the second set of simulations, p and r vary in such a way that σ_k increases and τ remains fixed. Because $k_y = k_z$, the first two simulations only investigate the estimation of Z . Analogous results would apply to Y .

Theorem 4.1 says, perhaps surprisingly, that the estimation performance only depends on n, P, τ , and σ_k . In the four parameter model, these quantities do not depend on Y . There are three setups corresponding to three different designs for Y and Z . In all designs, $k = 30$ and $s = 30$.

- In the second setup, $Y = Z$.
- In the first setup, Y and Z are chosen uniformly at random from the set of all possible partitions that ensure each block contains $s = 30$ nodes.
- In the third setup, the blocks are “non-overlapping”, meaning that if $y_i = y_j$, then $z_i \neq z_j$ for all pairs i and j .

These three designs range from Y and Z being the same to Y and Z being completely different. Simulations 1 and 2 investigate whether making Y and Z more similar improves the estimation of the partition. By examining these three different setups, the first two simulations show that the estimation of Z is neither improved nor hampered if Y and Z are more similar.

5.1.1. Simulation 1. One drawback of Theorem 4.1 is that it requires an asymptotically dense graph. Very few empirical networks display such an edge density. This simulation investigates the sensitivity of DI-SIM to a diminishing number of edges. To ensure that these results are not confounded by the effects of σ_k (the spectral gap), the values of p

and r change such that τ decreases while σ_k stays constant. By letting r vary and defining $p = 8r$, τ changes while

$$\sigma_k = \frac{1}{(k(r/p) + 1)^2} = \frac{1}{(30(1/8) + 1)^2} \approx .044.$$

Figure 1 displays the simulation results for a sequence of 100 equally spaced values of r between .02 and .05. To decrease the variability, each simulation was run 10 times. Only the average is displayed.




FIGURE 1. This simulation uses the four parameter Stochastic Co-Blockmodel with $k = 30$ and $s = 30$. The probabilities p and r vary such that $p = 8r$. This simulation shows that for very small values of τ , DI-SIM performs poorly. The three separate lines correspond to the three different designs for Y and Z described in the bullet points in Section 5.1. These lines do not appear significantly different suggesting that the estimation of Z is neither hampered nor improved if Y is made more similar to Z . (See the bullet points at the beginning of Section 5.1 for the three different designs of Y and Z .)

Figure 1 demonstrates two things. First, the number of misclustered nodes increases as $\tau \rightarrow 0$. This shows that DI-SIM performs poorly for small values of τ . Second, all three lines are nearly overlapping. This shows that the design of Y and Z do not affect

the performance of the algorithm. The two extremes are that $Y = Z$ or that the blocks are “non-overlapping,” $y_i = y_j \implies z_i \neq z_j$. The third line represents the random design, where block memberships are assigned randomly. DI-SIM is unaffected by these design choices.

5.1.2. Simulation 2. Theorem 4.1 requires that $n^{-1/2}(\log n)^2 = O(\sigma_k)$. This condition implies that the singular vectors that represent the block memberships are well separated from the singular vectors that do not represent the block memberships. This simulation investigates the sensitivity of DI-SIM to σ_k under the four parameter Stochastic Co-Blockmodel. To ensure that these simulation results are not confounded by the effects of τ , the values of p and r change such that τ remains constant. By letting p change and defining

$$r = \frac{1}{27} - \frac{p}{30},$$

τ remains fixed at $1/27 \approx .037$. Figure 2 displays the simulation results for a sequence of 100 equally spaced values of p between .05 and .36. To decrease the variability, each simulation was run 10 times. Only the average is displayed.

Figure 2 demonstrates two things. First, the lines are decreasing. This indicates that, as σ_k (the spectral gap) increases, DI-SIM can correctly cluster more nodes. Secondly, all three lines are nearly overlapping. This suggests that, under the four parameter Stochastic Blockmodel, the design of Y and Z do not interact with the spectral gap (σ_k) to impact the performance of the algorithm.

5.2. Simulation 3. In the two previous simulations, $k_y = k_z$. This simulation examines the unbalanced scenario, where k_y is fixed and k_z grows. In this simulation, each block in Y has an equal number of nodes (n/k_y) and each block in Z has an equal number of nodes (n/k_z). The size of each Z block is fixed at 20 ($20 = n/k_z$). Because k_z is growing, n is growing, and the size of each block in Y is also growing.

In the previous simulations, the matrix B is a diagonal matrix plus a constant matrix. In the unbalanced setting B is rectangular, so the previous model is no longer applicable. Instead, $B \in [0, 1]^{k_y \times k_z}$ is generated with iid $Uniform(.04, .4)$ random variables, removing columns until the columns are sufficiently separated. Specifically, the following pseudocode explains how B is simulated. Define B_i to be the i th column of B .

Sampling $B \in [0, 1]^{k_y \times k_z}$ for Simulation 3

Initialize $B_{uv} \sim Uniform(.04, .4)$ iid for all $u \in 1, \dots, k_y, v \in 1, \dots, k_z$.

for(v in $2 : k_z$)

 while $\min_{j=1, \dots, i-1} \|B_j - B_v\|_2 > .06k_y$

$B_{iv} \sim Uniform(.04, .4)$ iid for all $i \in 1, \dots, k_y$.




FIGURE 2. This simulation uses the four parameter Stochastic Co-Blockmodel with $k = 30$ and $s = 30$. The probabilities p and r vary such that $r = 1/27 - p/30$. This simulation shows that after σ crosses a threshold, between .02 and .04, DI-SIM performs well. The three separate lines correspond to the three different designs for Y and Z described in the bullet points in Section 5.1. These lines are indistinguishable, again suggesting that the estimation of Z is neither harder nor easier if Y is made more similar to Z . (See the bullet points at the beginning of Section 5.1 for the three different designs of Y and Z .)

The constraint that $\|B_j - B_v\|_2 > .06k_y$ helps prevent γ , defined in Equation (4.3), from becoming too small. The definition of γ relies on $B^{\mathcal{L}}$. However, for computational considerations, the constraint $\|B_j - B_v\|_2 > .06k_y$ is enforced on B and not $B^{\mathcal{L}}$.

Figure 3 suggests that the proportion of misclustered nodes converges to zero faster than Theorem 4.1 implies. The horizontal axis in this figure corresponds to k_z growing from 5 to 65 in intervals of 5. Each Z block contains 20 nodes. Each Y block has an equal number of nodes and this size grows with k_z . At each point, B is generated 50 times. For each of these B 's, one graph is simulated. The dashed line corresponds to the theoretical rate of convergence (Equation 4.5 divided by n) averaged over the 50 simulations. The solid line corresponds to the proportion of nodes that DI-SIM Z -misclusters averaged over the

50 simulations. The slope of the solid line appears steeper than the dashed line, suggesting that there is room to improve our theoretical results in Theorem 4.1.

Figure 3 does not display the average number of Y -misclustered nodes. For every simulation in which $k_z > 10$, there were no Y -misclustered nodes, $|\mathcal{M}_y| = 0$.




FIGURE 3. This simulation sets $k_y = 5$ and allows k_z to grow from 5 to 65 in intervals of 5. At each point, the data is generated 50 times. The dashed line represents the theoretical rate of convergence for the proportion of misclustered nodes (this is Equation (4.5) divided by n). The solid line gives the proportion of nodes Z -misclustered. Notice that the solid line appears to have a faster slope than the dashed line. This suggests there is room for improvement in our theoretical results.

6. DISCUSSION

By extending both spectral clustering and the Stochastic Blockmodel to a co-clustering framework, this paper aims to better conceptualize clustering in directed graphs. Our hope is to present co-clustering as a meaningful procedure for directed networks and that this helps to guide the development of reasonable questions and sensible similarity measures for network researchers.

In particular, Section 2 introduces the DI-SIM algorithm. By using the singular value decomposition, this algorithm uses the similarity measures “number of common parents” and “number of common offspring” to co-cluster the nodes into two different partitions. Section 3 motivates and introduces the Stochastic Co-Blockmodel that encodes the concepts of co-clustering in a statistical model. The classical Stochastic Blockmodel employs the

fundamental concept of stochastic equivalence; two nodes in the same block are stochastically equivalent. Section 3 argues that one can allow for two distinct concepts of stochastic equivalence in directed graphs,

Stochastically equivalent senders: $P(a \rightarrow x) = P(b \rightarrow x) \forall x$

Stochastically equivalent receivers: $P(x \rightarrow a) = P(x \rightarrow b) \forall x$.

This prompts a new parameterization of the Stochastic Blockmodel for directed graphs. This new model contains two partitions of the nodes. The first partition corresponds to the first form of stochastic equivalence above. The second partition corresponds to the second form of stochastic equivalence. In the classical Stochastic Blockmodel, where there is one type of stochastic equivalence, the estimation of the blocks is equivalent to clustering. In this new model, with two types of stochastic equivalence, the estimation of the two types of blocks is equivalent to co-clustering.

Theorem 4.1 in Section 4, this paper’s main result, shows that the clusters from DI-SIM estimate the two different partitions in the Stochastic Co-Blockmodel, even in the high-dimensional setting where the number of blocks increases with the number of nodes. In other words, under certain conditions, the two sets of clusters from DI-SIM estimate the two different types of stochastic equivalence encoded in the Stochastic Co-Blockmodel. Asymptotically in the number of nodes, Theorem 4.1 bounds the number of “misclustered” nodes as long as (1) the spectral gap is not too small and (2) the minimum expected degree grows fast enough. These results are analogous to the results on spectral clustering under the Stochastic Blockmodel Rohe et al. [2011]. The simulations in Section 5 demonstrate that breaking these conditions can drastically diminish the ability of DI-SIM to estimate the blocks in the Stochastic Co-Blockmodel.

The main limitation of Theorem 4.1 is that it does not apply to sparse graphs. Rather, this theorem requires that the minimum expected degree grows at the same rate as the number of nodes (ignoring $\log n$ terms). In large empirical networks, edges are not dense enough to suggest this type of asymptotic framework. One area for future research is to study the statistical properties of spectral clustering in a sparse graph asymptotic framework. Previous clustering results for the Stochastic Blockmodel suggest that it is possible to cluster the nodes when the graph is sparse; Bickel and Chen [2009] studied the MLE and modularity methods under the low dimensional Stochastic Blockmodel, Choi et al. [2012] studied the MLE under a high dimensional Stochastic Blockmodel where the number of clusters grows like $n^{1/2}$, and Zhao et al. [2011] studied the MLE and modularity methods under the low dimensional, degree-corrected Stochastic Blockmodel. In all of these papers, the expected degree can grow like $\log^{3+\epsilon} n$. Unfortunately, the clusters estimated by maximum likelihood and modularity techniques are computationally intensive, and potentially NP hard, to compute. Future research should investigate the performance of spectral techniques in an asymptotic setting that allows for sparse connections.

APPENDIX A. DIRECTED LATENT SPACE MODEL

The following definition of the directed latent space model is motivated by the Aldous-Hoover representation for infinite exchangeable arrays and the latent space model proposed by Hoff et al. [2002]. It specifies the distribution of the random directed adjacency matrix $A \in \{0, 1\}^{n \times n}$.

Definition 5. *The random adjacency matrix A is from the **directed latent space model** if and only if*

$$\mathbb{P}(A | \{z_i, y_i\}_{i=1}^n) = \prod_{i < j} \mathbb{P}(A_{ij} | y_i, z_j)$$

where $\{z_i, y_i\}_{i=1}^n \subset \mathbb{R}^k \times \mathbb{R}^k$ are pairs of random vectors that are independent across $i = 1, \dots, n$.

In this definition, $\mathbb{P}(A_{ij} | y_i, z_j)$ is the probability mass function of A_{ij} conditioned on y_i and z_j . Define $Y \in \mathbb{R}^{n \times k}$ such that its i th row is y_i for all $i \in V$. Similarly, define $Z \in \mathbb{R}^{n \times k}$ such that its i th row is z_i . **Throughout this paper we condition on Y and Z .** Because $\mathbb{P}(A_{ij} = 1 | Y, Z) = \mathbb{E}(A_{ij} | Y, Z)$, the model is then completely parametrized by the matrix

$$\mathcal{A} = \mathbb{E}(A | Y, Z) \in \mathbb{R}^{n \times n},$$

where \mathcal{A} depends on Y and Z , but this is dropped for notational convenience.

The Stochastic Blockmodel, introduced by Holland et al. [1983], is a specific latent space model with well defined communities. The following definition extends the Stochastic Blockmodel to allow for the asymmetric communities discussed in the previous section.

Definition 6. *The **Stochastic Co-Blockmodel** with k blocks is a directed latent space model with*

$$\mathcal{A} = Y B Z^T,$$

where $Y, Z \in \{0, 1\}^{n \times k}$ both have exactly one 1 in each row and at least one 1 in each column and $B \in [0, 1]^{k \times k}$ is full rank.

APPENDIX B. CONVERGENCE OF SINGULAR VECTORS

The classical spectral clustering algorithm above can be divided into two steps; (1) find the eigendecomposition of $L^{(s)}$ and (2) run k -means. Rohe et al. [2011] studied the estimation performance of the classical spectral clustering algorithm under a standard social network model. In this analysis, standard perturbation results were not enough to control the eigenvectors of the random matrix $L^{(s)}$. Instead, the paper devised the “squaring trick.” This paper also utilizes the squaring trick.

This trick is the composition of two observations. The first observation is easily demonstrated on the adjacency matrix of a directed Erdos-Renyi random graph. Under this

model, each element of $A \in \{0, 1\}^{n \times n}$ is a Bernoulli random variable with probability p . As a result, A_{ij} is either a zero or a one, independent of the size of the matrix n . However, the elements of AA are equal to the sum of n zero-one variables,

$$[AA]_{ij} = \sum_{\ell} A_{i\ell} A_{\ell j}.$$

By applying concentration results to this sum of (almost) independent random variables, standard perturbation theorems imply that the eigenvectors of AA are close to the eigenvectors of $E(A)E(A)$. The second observation in the squaring trick is the following:

$$Ax = \lambda x \Rightarrow AAx = A(\lambda x) = \lambda^2 x.$$

This shows that any eigenvector of A is also an eigenvector of AA . Taken together, these observations show that the eigenvectors of A are close to the eigenvectors of $E(A)$. These results are easiest to state for the matrix A under the Erdos-Renyi random graph model. However, analogous results hold for the graph Laplacian $L^{(s)}$ under the more general latent space model [Hoff et al., 2002, Rohe et al., 2011].

To study the singular vectors of L , this appendix first studies the convergence of $L^T L$ under the directed latent space model. The results in this paper are asymptotic in the number of nodes n . When it is appropriate, the matrices above are given a superscript of n to emphasize this dependence. Other times, this superscript is discarded for notational convenience.

Recall that

$$(B.1) \quad \tau_n = \min_{i=1, \dots, n} \min\{\mathcal{P}_{ii}^{(n)}, \mathcal{O}_{ii}^{(n)}\}/n.$$

Because $\mathcal{P}_{ii}^{(n)}$ is the expected out-degree for node i and $\mathcal{O}_{ii}^{(n)}$ is the expected in-degree for node i , τ_n is the minimum expected degree divided by the maximum possible degree. It measures how quickly the number of edges accumulates.

Theorem B.1. *Define the sequence of random matrices $A^{(n)} \in \{0, 1\}^{n \times n}$ to be from a sequence of directed latent space models with population matrices $\mathcal{A}^{(n)} \in [0, 1]^{n \times n}$. With $A^{(n)}$, define the observed graph Laplacian $L^{(n)}$ as in (2.3). Let $\mathcal{L}^{(n)}$ be the population version of $L^{(n)}$ as defined in Equation (C.1). Define τ_n as in Equation (B.1).*

If there exists $N > 0$, such that $\tau^2 \log(n) > 2$ for all $n > N$, then

$$\|(L^{(n)})^T L^{(n)} - (\mathcal{L}^{(n)})^T \mathcal{L}^{(n)}\|_F = o\left(\frac{\log(n)}{\tau_n^2 n^{1/2}}\right) \quad a.s.$$

The proof of Theorem B.1 relies on the following lemma.

Lemma B.1. *Under the directed latent space mode, if $n^{1/2}/\log(n) > 2$, then*

$$\mathbb{P}\left(\|L^T L - \mathcal{L}^T \mathcal{L}\|_F \geq \frac{32 \log(n)}{\tau^2 n^{1/2}}\right) \leq 6n^{2-2\tau^2 \log n}.$$

The same statement holds for $\|LL^T - \mathcal{L}\mathcal{L}^T\|_F$.

Proof. The main complication of proving Lemma B.1 is controlling the dependencies between the elements of $L^T L$. We do this with an intermediate step that uses the matrix

$$\tilde{L} = \mathcal{O}^{-1/2} A \mathcal{P}^{-1/2}$$

and two sets Γ and Λ . Γ constrains the matrices P and O , while Λ constrains the matrix $A\mathcal{O}^{-1}A$. These sets will be defined in the proof. To ease the notation, define

$$\mathbb{P}_{\Gamma\Lambda}(B) = \mathbb{P}(B \cap \Gamma \cap \Lambda)$$

where B is some event.

This proof shows that under the sets Γ and Λ the probability of the norm exceeding $32 \log(n) \tau^{-2} n^{-1/2}$ is exactly zero for large enough n and that the probability of Γ or Λ converges to one. To ease notation, define $a = 32 \log(n) \tau^{-2} n^{-1/2}$.

$$\mathbb{P}(\|L^T L - \mathcal{L}^T \mathcal{L}\|_F \geq a) \leq \mathbb{P}_{\Gamma\Lambda} \left(\sum_{i,j} [L^T L - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2 \right) + \mathbb{P}((\Gamma \cap \Lambda)^c)$$

Decompose this first quantity with the union bound.

$$\begin{aligned} \mathbb{P}_{\Gamma\Lambda} \left(\sum_{i,j} [L^T L - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2 \right) &\leq \sum_{i,j} \mathbb{P}_{\Gamma\Lambda} ([L^T L - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2/n^2) \\ \text{(B.2)} \quad &\leq \sum_{i,j} \mathbb{P}_{\Gamma\Lambda} ([L^T L - \tilde{L}^T \tilde{L}]_{ij}^2 \geq a^2/2n^2) \\ &\quad + \mathbb{P}_{\Gamma\Lambda} ([\tilde{L}^T \tilde{L} - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2/2n^2) \end{aligned}$$

To constrain the terms $[\tilde{L}^T \tilde{L} - \mathcal{L}^T \mathcal{L}]_{ij}^2$, define

$$\Lambda = \bigcap_{i,j} \left\{ \left| \sum_{k=1}^n (A_{ki} A_{kj} - \mathcal{A}_{ki} \mathcal{A}_{kj}) / \mathcal{O}_{kk} \right| < n^{-1/2} \log n \right\}.$$

Under Λ ,

$$|\tilde{L}^T \tilde{L} - \mathcal{L}^T \mathcal{L}|_{ij} = \frac{1}{\sqrt{\mathcal{P}_{ii} \mathcal{P}_{jj}}} \left| \sum_{k=1}^n (A_{ki} A_{kj} - \mathcal{A}_{ki} \mathcal{A}_{kj}) / \mathcal{O}_{kk} \right| < \frac{1}{\tau n} \frac{\log n}{\sqrt{n}} \leq \frac{a}{2n}$$

This controls the second set of terms in line B.2. The next lines show that

$$\mathbb{P}_{\Gamma\Lambda} ([L^T L - \tilde{L}^T \tilde{L}]_{ij}^2 \geq a^2/2n^2) = 0.$$

For $b(n) = \log(n)n^{-1/2}$ define the following set.

$$\Gamma = \bigcap_i \left\{ O_{ii} \in \mathcal{O}_{ii} [1 - b(n), 1 + b(n)] \ \& \ C_{ii} \in \mathcal{P}_{ii} [1 - b(n), 1 + b(n)] \right\}$$

Define another set,

$$\Gamma' = \bigcap_{i,j,k} \left\{ \frac{1}{C_{kk}(O_{ii}O_{jj})^{1/2}} \in \frac{[1 - 16b(n), 1 + 16b(n)]}{\mathcal{P}_{kk}(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} \right\}.$$

Lemma B.2. *Under the definitions and assumptions for the current proof, $\Gamma \subset \Gamma'$.*

A proof of this lemma follows the current proof.

Under the set Γ , and thus Γ' ,

$$\begin{aligned} |L^T L - \tilde{L}^T \tilde{L}|_{ij} &= \left| \sum_k \left(\frac{A_{ik}A_{kj}}{C_{kk}(O_{ii}O_{jj})^{1/2}} - \frac{A_{ik}A_{kj}}{\mathcal{P}_{kk}(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} \right) \right| \\ &\leq \sum_k \left| \frac{1}{C_{kk}(O_{ii}O_{jj})^{1/2}} - \frac{1}{\mathcal{P}_{kk}(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} \right| \\ &\leq \sum_k \left| \frac{16b(n)}{\mathcal{P}_{kk}(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} \right| \\ &\leq \sum_k \frac{16b(n)}{\tau^2 n^2} \\ &\leq \frac{16b(n)}{\tau^2 n}. \\ &\leq \frac{16 \log n}{\tau^2 n^{3/2}}. \\ &= \frac{a}{2n}. \end{aligned}$$

Using line B.2, this shows that

$$\mathbb{P}_{\Gamma\Lambda} \left(\sum_{i,j} [L^T L - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2 \right) = 0.$$

The remaining step is to bound

$$\mathbb{P}((\Gamma \cap \Lambda)^c) \leq \mathbb{P}(\Gamma^c) + \mathbb{P}(\Lambda^c).$$

Using the union bound Hoeffding bound,

$$\begin{aligned}
\mathbb{P}(\Gamma^c) &\leq \sum_i \mathbb{P}(O_{ii} \notin \mathcal{O}_{ii}[1 - b(n), 1 + b(n)]) + \mathbb{P}(C_{ii} \notin \mathcal{P}_{ii}[1 - b(n), 1 + b(n)]) \\
&< \sum_i 2 \exp \left(-2 \left(\frac{\mathcal{O}_{ii} \log n}{\sqrt{n}} \right)^2 \frac{1}{n} \right) + 2 \exp \left(-2 \left(\frac{\mathcal{P}_{ii} \log n}{\sqrt{n}} \right)^2 \frac{1}{n} \right) \\
&\leq 4n \exp(-2\tau^2 (\log n)^2) \\
&= 4n^{1-2\tau^2 \log n}.
\end{aligned}$$

Again, using the union bound and the Hoeffding bound,

$$\begin{aligned}
\mathbb{P}(\Lambda^c) &\leq \sum_{i,j} \mathbb{P} \left(\left| \sum_k (A_{ki} A_{kj} - \mathcal{A}_{ki} \mathcal{A}_{kj}) / \mathcal{O}_{kk} \right| > n^{-1/2} \log n \right) \\
&< \sum_{i,j} 2 \exp \left(\frac{-2(\log n)^2}{n} \frac{1}{\sum_k 1/\mathcal{O}_{kk}} \right) \\
&\leq \sum_{i,j} 2 \exp(-2(\log n)^2 \tau^2) \\
&\leq 2n^2 \exp(-2(\log n)^2 \tau^2) \\
&\leq 2n^{2-2\tau^2 \log n}.
\end{aligned}$$

Putting the pieces together,

$$\begin{aligned}
\mathbb{P}(\|L^T L - \mathcal{L}^T \mathcal{L}\|_F \geq a) &\leq \mathbb{P}_{\Gamma\Lambda} \left(\sum_{i,j} [L^T L - \mathcal{L}^T \mathcal{L}]_{ij}^2 \geq a^2 \right) + \mathbb{P}(\Gamma^c) + \mathbb{P}(\Lambda^c) \\
&< 0 + 4n^{1-2\tau^2 \log n} + 2n^{2-2\tau^2 \log n} \\
&\leq 6n^{2-2\tau^2 \log n}.
\end{aligned}$$

After proving Lemma B.2, the proof of Lemma B.1 is complete. \square

Proof. This proves Lemma B.2. To simplify the notation, define $u(n) = 1 + b(n)$, $l(n) = 1 - b(n)$. Define the sets

$$\begin{aligned}\Gamma(1) &= \bigcap_i \left\{ \frac{1}{C_{ii}} \in \frac{1}{\mathcal{P}_{ii}} [u(n)^{-1}, l(n)^{-1}] \right\} \\ \Gamma(2) &= \bigcap_{i,j} \left\{ \frac{1}{(O_{ii}O_{jj})^{1/2}} \in \frac{1}{(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} [u(n)^{-1}, l(n)^{-1}] \right\} \\ \Gamma(3) &= \bigcap_{i,j,k} \left\{ \frac{1}{C_{kk}(O_{ii}O_{jj})^{1/2}} \in \frac{[u(n)^{-2}, l(n)^{-2}]}{\mathcal{P}_{kk}(\mathcal{O}_{ii}\mathcal{O}_{jj})^{1/2}} \right\}\end{aligned}$$

Notice that $\Gamma \subseteq \{\Gamma(1) \cup \Gamma(2)\} \subseteq \Gamma(3)$. It is sufficient to show $\Gamma(3) \subset \Gamma'$. This is true because

$$\begin{aligned}\frac{1}{u(n)^2} = \frac{1}{(1+b(n))^2} &= \frac{b(n)^{-2}}{(b(n)^{-1}+1)^2} > \frac{b(n)^{-2}-1}{(b(n)^{-1}+1)^2} = \frac{(b(n)^{-1}-1)(b(n)^{-1}+1)}{(b(n)^{-1}+1)^2} \\ &= \frac{b(n)^{-1}-1}{b(n)^{-1}+1} = 1 - \frac{2}{b(n)^{-1}+1} > 1 - 16b(n).\end{aligned}$$

The 16 in the last bound is larger than it needs to be so that the upper and lower bounds in Γ' are symmetric. For the other direction,

$$\begin{aligned}\frac{1}{l(n)^2} = \frac{1}{(1-b(n))^2} &= \frac{b(n)^{-2}}{(b(n)^{-1}-1)^2} = \left(1 + \frac{1}{b(n)^{-1}-1}\right)^2 \\ &= 1 + \frac{2}{b(n)^{-1}-1} + \frac{1}{(b(n)^{-1}-1)^2}.\end{aligned}$$

To bound the last two elements, recall that it is assumed $\sqrt{n}/\log(n) > 2$. Equivalently, $1 - b(n) > 1/2$. This yields both of the following:

$$\frac{1}{(b(n)^{-1}-1)^2} < \frac{2}{b(n)^{-1}-1} \quad \text{and} \quad \frac{2}{b(n)^{-1}-1} = \frac{2b(n)}{1-b(n)} < 8b(n).$$

Putting these together,

$$\frac{1}{l(n)^2} < 1 + 16b(n).$$

This shows that $\Gamma \subset \Gamma'$. □

The following proves Theorem B.1.

Proof. Adding the n super- and subscripts to Lemma B.1, it states that if $n^{1/2}/\log(n) > 2$, then

$$\mathbb{P} \left(\|(L^{(n)})^T L^{(n)} - (\mathcal{L}^{(n)})^T \mathcal{L}^{(n)}\|_F \geq \frac{32 \log(n)}{\tau^2 n^{1/2}} \right) \leq 6n^{2-2\tau^2 \log n}.$$

By assumption, for all $n > N$, $\tau^2 \log(n) > 2$. This implies that $2 - 2\tau^2 \log(n) < -2$ for all $n > N$. Rearranging and summing over n , for any fixed $\epsilon > 0$,

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbb{P} \left(\frac{\|L^{(n)} L^{(n)} - \mathcal{L}^{(n)} \mathcal{L}^{(n)}\|_F}{32\tau_n^{-2} \log(n) n^{-1/2}/\epsilon} \geq \epsilon \right) &\leq N + 6 \sum_{n=N+1}^{\infty} n^{2-2\tau^2 \log(n)} \\ &\leq N + 6 \sum_{n=N+1}^{\infty} n^{-2}, \end{aligned}$$

which is a summable sequence. By the Borel-Cantelli Theorem,

$$\|(L^{(n)})^T L^{(n)} - (\mathcal{L}^{(n)})^T \mathcal{L}^{(n)}\|_F = o(\tau_n^{-2} \log(n) n^{-1/2}) \quad a.s.$$

□

In order to prove Theorem B.2, we need the following version of the Davis-Kahan Theorem.

Proposition B.1. (Davis-Kahan) *Let $S \subset \mathbb{R}$ be an interval. Denote X as an orthonormal matrix whose column space is equal to the eigenspace of $A = A^T$ corresponding to the eigenvalues of A contained in S (more formally, the column space of X is the image of the spectral projection of A induced by S). Denote by \tilde{X} the analogous quantity for $\tilde{A} = \tilde{A}^T$. Define the distance between S and the spectrum of A outside of S as*

$$\delta = \min\{|\ell - s|; \ell \text{ eigenvalue of } A, \ell \notin S, s \in S\}.$$

If X and \tilde{X} are of the same dimension, then

$$\frac{1}{2} \|\tilde{X} - X\mathcal{R}\|_F^2 \leq \frac{\|\tilde{A} - A\|_F^2}{\delta^2}$$

where \mathcal{R} is an orthonormal matrix. With singular value decomposition, define $X^T \tilde{X} = U\Sigma V^T$. Then, $\mathcal{R} = UV^T$.

The original Davis-Kahan Theorem bounds the “canonical angle,” also known as the “principal angle,” between the column spaces of X and \tilde{X} . The appendix in Rohe et al. [2011] explains how the original theorem can be converted into Proposition B.1.

Proposition B.1 aids the following proof of Theorem B.2.

Proof. To show that $k_n = \mathcal{K}_n$ eventually, define $\sigma_{n1} \geq \dots \geq \sigma_{nn}$ to be the singular values of $L^{(n)}$ and define $\sigma_{n1}^* \geq \dots \geq \sigma_{nn}^*$ to be the singular values of $\mathcal{L}^{(n)}$.

$$\begin{aligned} \max_i |\sigma_{ni} - \sigma_{ni}^*| &\leq \|(L^{(n)})^T L^{(n)} - (\mathcal{L}^{(n)})^T \mathcal{L}^{(n)}\|_F \\ (B.3) \quad &= o\left(\frac{\log(n)}{\tau_n^2 n^{1/2}}\right) \quad a.s. \end{aligned}$$

This follows from Weyl's inequality [Bhatia, 2007], the fact that the Frobenius norm is an upper bound of the spectral norm, and Theorem B.1. By assumption,

$$\frac{\log(n)}{\tau_n^2 n^{1/2}} \leq \frac{(\log n)^2}{2 n^{1/2}} = O(\min\{\delta_n, \delta'_n\}).$$

This shows that

$$\max_i |\sigma_{ni} - \sigma_{ni}^*| = o(\min\{\delta_n, \delta'_n\}).$$

By the definition of δ_n and δ'_n , it follows that eventually, for $i = 1, \dots, n$,

$$\sigma_{ni} \in S_n \Leftrightarrow \sigma_{ni}^* \in S_n.$$

Therefore, $k_n = \mathcal{K}_n$ eventually.

When $k_n = \mathcal{K}_n$, the results follow from Proposition B.1. To define the matrix \mathcal{R}_n , use singular value decomposition to define $(\mathcal{V}_n)^T V_n = M_{n1} \Sigma_n M_{n2}^T$. Then, $\mathcal{R}_n = M_{n1} M_{n2}^T$. \square

The next theorem uses the above lemma in concert with the Davis-Kahan Theorem to show that the left and right singular vectors of $L^{(n)}$ converge to the left and right singular vectors of $\mathcal{L}^{(n)}$.

Define the open interval $S_n \subset \mathbb{R}$ to contain the squared singular values of $\mathcal{L}^{(n)}$ that are of interest. Define

$$(B.4) \quad \delta_n = \min\{|\sigma^2 - s|; \sigma \text{ is a singular value of } \mathcal{L}^{(n)}, \sigma \notin S, s \in S\}$$

$$(B.5) \quad \delta'_n = \inf\{|\sigma^2 - s|; \sigma \text{ is a singular value of } \mathcal{L}^{(n)}, s \notin S_n\}$$

The quantity δ_n measures the distance between S_n and the singular values that are not of interest. If δ_n is too small, then $L^{(n)}$ might have too many singular values inside S_n . The quantity δ'_n measures how well S_n insulates the singular values of interest. If δ'_n is too small, then some important singular values in $L^{(n)}$ might fall outside of S_n . By restricting the rate at which δ_n and δ'_n converge to zero, the next theorem ensures that the ‘‘eigengap’’ is not too small.

Theorem B.2. Define $A^{(n)} \in \{0, 1\}^{n \times n}$ to be a sequence of growing random adjacency matrices from the directed latent space model with population matrices $\mathcal{A}^{(n)}$. Define the observed graph Laplacian $L^{(n)}$ as in (2.3). Let $\mathcal{L}^{(n)}$ be the population version of $L^{(n)}$ as defined in Equation (C.1).

Define $S_n \subset \mathbb{R}$ as a sequence of intervals. Let k_n denote the number of squared singular values of $L^{(n)}$ that fall within S_n . Let \mathcal{K}_n denote the number of squared singular values of $\mathcal{L}^{(n)}$ that fall within S_n .

Let $V_n \in \mathbb{R}^{n \times k_n}$ be an orthonormal matrix where each column is a right singular vector of $L^{(n)}$ whose corresponding squared singular value falls within S_n . Let $\mathcal{V}_n \in \mathbb{R}^{n \times \mathcal{K}_n}$ be an orthonormal matrix where each column is a right singular vector of $\mathcal{L}^{(n)}$ whose corresponding squared singular value falls within S_n .

Assume there exists $N \in \mathbb{R}$ such that for all $n > N$, $\tau_n^2 > 2/\log n$. Define δ_n and δ'_n as defined in Equations B.4 and B.5. If, $n^{-1/2}(\log n)^2 = O(\min\{\delta_n, \delta'_n\})$, then eventually $k_n = \mathcal{K}_n$. After that point,

$$\|V_n - \mathcal{V}_n \mathcal{R}_n\|_F = o\left(\frac{\log n}{\delta_n \tau_n^2 n^{1/2}}\right) a.s.$$

where \mathcal{R}_n is an orthonormal rotation that depends on V_n and \mathcal{V}_n .

APPENDIX C. CLUSTERING

To rigorously discuss the asymptotic estimation properties of DI-SIM, the next subsections (1) define a population version of the graph Laplacian, (2) examine the behavior of DI-SIM applied to a population version of the graph Laplacian, and (3) compare this to DI-SIM applied to the observed graph Laplacian.

To give this definition, the next two subsections study DI-SIM applied to \mathcal{L} and compare these results to DI-SIM applied to L . This discussion largely follows the discussion in Rohe et al. [2011].

C.1. The population version of DI-SIM. Define $\mathcal{A} = E(A)$ as the population version of the adjacency matrix A . Recall that under the Stochastic Co-Blockmodel,

$$\mathcal{A} = YBZ^T,$$

where $Y \in \{0, 1\}^{n \times k_y}$, $Z \in \{0, 1\}^{n \times k_z}$, and $B \in [0, 1]^{k_y \times k_z}$. All of Section 4 will assume that $k_y \leq k_z$, without loss of generality.

Define population versions of O , P , and L all in $\mathbb{R}^{n \times n}$ as

$$\begin{aligned} \mathcal{O}_{jj} &= \sum_k \mathcal{A}_{kj} \\ \mathcal{P}_{ii} &= \sum_k \mathcal{A}_{ik} \\ \mathcal{L} &= \mathcal{P}^{-1/2} \mathcal{A} \mathcal{O}^{-1/2} \end{aligned} \tag{C.1}$$

where \mathcal{O} and \mathcal{P} are diagonal matrices.

This subsection shows that DI-SIM applied to \mathcal{L} can perfectly identify the blocks in the Stochastic Co-Blockmodel.

To determine the clusters based on common parents, recall that DI-SIM applied to L ,

- (1) finds the right singular vectors $U \in \mathbb{R}^{n \times k_y}$,
- (2) defines the n rows of U as $u_1, \dots, u_n \in \mathbb{R}^{k_y}$,
- (3) runs k -means on u_1, \dots, u_n with k_y clusters.
- (4) repeats (1) for the the left singular vectors $V \in \mathbb{R}^{n \times k_y}$ with k_z clusters.

This statement of the algorithm uses the assumption that $k_y \leq k_z$.

k -means clusters points u_1, \dots, u_n in euclidean space by optimizing the following objective function [Steinhaus, 1956],

$$(C.2) \quad \min_{\{m_1, \dots, m_{k_y}\} \subset \mathbb{R}^{k_y}} \sum_i \min_g \|u_i - m_g\|_2^2.$$

Define the *centroids* as the arguments $m_1^*, \dots, m_{k_y}^*$ that optimize (C.2). The analysis in this paper addresses the true optimum of (C.2).

To examine the results of DI-SIM applied to \mathcal{L} , define

$$\begin{aligned} D_y &= \text{diag}(BZ^T \mathbf{1}_n) \in \mathbb{R}^{k_y \times k_y} \\ D_z &= \text{diag}(\mathbf{1}_n^T YB) \in \mathbb{R}^{k_z \times k_z} \\ B^{\mathcal{L}} &= D_y^{-1/2} B D_z^{-1/2} \in [0, 1]^{k_y \times k_z}, \end{aligned}$$

where $\text{diag}(x)$ for $x \in \mathbb{R}^d$ is a diagonal matrix in $\mathbb{R}^{d \times d}$ with $\text{diag}(x)_{jj} = x_j$. It follows that there is an alternative expression for \mathcal{L} ,

$$\mathcal{L} = Y B^{\mathcal{L}} Z^T.$$

By the singular value decomposition, there exist orthonormal matrices $\mathcal{U}, \mathcal{V} \in \mathbb{R}^{n \times k_y}$ and diagonal matrix $\Lambda \in \mathbb{R}^{k_y \times k_y}$ such that

$$(C.3) \quad \mathcal{L} = \mathcal{U} \Lambda \mathcal{V}^T.$$

The next lemma shows that DI-SIM applied to the population Laplacian, \mathcal{L} , can discover the block structure in the matrices Y and Z . Lemma C.1 is essential to defining “misclustered.”

Lemma C.1. *Under the Stochastic Co-Blockmodel with $k_y \leq k_z$, define \mathcal{L} as in Equation (C.1) and define \mathcal{U}, \mathcal{V} , and Λ as in Equation C.3.*

(1) *If $\text{rank}(B) = k_y$, then*

$$(C.4) \quad \mathcal{U}_i = \mathcal{U}_j \Leftrightarrow y_i = y_j,$$

where y_i is the i th row of Y and \mathcal{U}_i is the i th row of \mathcal{U} .

(2) *If $\text{rank}(B) = k_y$ and the columns of $B^{\mathcal{L}}$ are distinct, then*

$$(C.5) \quad \mathcal{V}_i = \mathcal{V}_j \Leftrightarrow z_i = z_j,$$

where z_i is the i th row of Z and \mathcal{V}_i is the i th row of \mathcal{V} .

This lemma implies that there exist matrices $\mu^y \in \mathbb{R}^{k_y \times k_y}$ and $\mu^z \in \mathbb{R}^{k_z \times k_y}$ such that $Y\mu^y = \mathcal{U}$ and $Z\mu^z = \mathcal{V}$.

Proof. There are two results in Lemma C.1. The first result concerns the left singular vectors and the second result concerns the right singular vectors. To simplify later results, the the proof of the first result is vastly different than the proof of the second result.

To prove that

$$\mathcal{U}_i = \mathcal{U}_j \Leftrightarrow y_i = y_j,$$

construct a matrix μ^y such that

$$\mathcal{U}_i = y_i \mu^y.$$

Recall that $\mathcal{L} = Y B^{\mathcal{L}} Z^T$. Define $B^{\mathcal{L}\mathcal{L}} = B^{\mathcal{L}} Z^T Z (B^{\mathcal{L}})^T$. Then,

$$\mathcal{L} \mathcal{L}^T = Y B^{\mathcal{L}\mathcal{L}} Y^T.$$

Because $B^{\mathcal{L}\mathcal{L}}$ is symmetric, so is $(Y^T Y)^{1/2} B^{\mathcal{L}\mathcal{L}} (Y^T Y)^{1/2}$. Further, by the assumptions of the Stochastic Co-Blockmodel and the assumption that $\text{rank}(B) = k_y$, $(Y^T Y)^{1/2} B^{\mathcal{L}\mathcal{L}} (Y^T Y)^{1/2}$ is full rank. By eigendecomposition, there exists an orthonormal matrix $U_{k_y} \in \mathbb{R}^{k_y \times k_y}$ that contains the eigenvectors in its columns and a diagonal matrix $\Lambda' \in \mathbb{R}^{k \times k}$ (with nonzero entries down the diagonal) such that

$$(Y^T Y)^{1/2} B^{\mathcal{L}\mathcal{L}} (Y^T Y)^{1/2} = U_{k_y} \Lambda' U_{k_y}^T.$$

Left multiply by $Y(Y^T Y)^{-1/2}$ and right multiply by $(Y^T Y)^{-1/2} Y^T$.

$$Y B^{\mathcal{L}\mathcal{L}} Y^T = (Y \mu^y) \Lambda' (Z \mu^y)^T,$$

for $\mu^y = (Y^T Y)^{-1/2} U_{k_y}$. Note that the left hand side of Equation C.1 is equal to $\mathcal{L} \mathcal{L}^T$. Note that $(Y \mu^y)^T Y \mu^y$ is equal to the identity. So, left multiplying Equation C.1 by $Y \mu^y$ shows that the columns of $Y \mu^y$ are the left singular vectors of \mathcal{L} and the diagonal of Λ' contains the squared singular values. Thus $\Lambda^2 = \Lambda'$,

$$\mathcal{U} = Y \mu^y, \text{ and } \mathcal{U}_i = y_i \mu^y.$$

Because $\det(\mu^y) = \det((Y^T Y)^{-1/2}) \det(U_k) > 0$, $\mu^y \in \mathbb{R}^{k_y \times k_y}$ is full rank. Therefore, $(\mu^y)^{-1}$ exists and

$$y_i \mu^y = y_j \mu^y \Leftrightarrow y_i = y_j.$$

The second part of Lemma C.1 says that, if $\text{rank}(B) = k_y$ and the columns of $B^{\mathcal{L}}$ are distinct, then $\mathcal{V}_i = \mathcal{V}_j \Leftrightarrow z_i = z_j$. Notice that

$$\mathcal{L} = Y B^{\mathcal{L}} Z^T = \mathcal{U} \Lambda \mathcal{V}^T = Y \mu^y \Lambda \mathcal{V}^T.$$

Left multiply by $\Lambda^{-1} (\mu^y)^{-1} (Y^T Y)^{-1} Y^T$ and take the transpose to get,

$$\mathcal{V} = Z (B^{\mathcal{L}})^T ((\mu^y)^{-1})^T \Lambda^{-1}.$$

Define $\mu^z = (B^{\mathcal{L}})^T ((\mu^y)^{-1})^T \Lambda^{-1}$. So, $z_i = z_j \Rightarrow v_i = v_j$. To prove the other direction, let $Z_{iu} = 1$ and $Z_{jv} = 1$ (this means that node i is in the u th block of Z and node j is in the v th block of Z) and define $B_{\cdot j}^{\mathcal{L}}$ as the j th column of $B^{\mathcal{L}}$.

$$\begin{aligned} \|\mathcal{V}_i - \mathcal{V}_j\|_2 = \|z_i \mu^z - z_j \mu^z\|_2 &= \|(B_{\cdot u}^{\mathcal{L}} - B_{\cdot v}^{\mathcal{L}})^T ((\mu^y)^{-1})^T \Lambda^{-1}\|_2 \\ &\geq \|B_{\cdot u}^{\mathcal{L}} - B_{\cdot v}^{\mathcal{L}}\|_2 \|(\mu^y)^{-1}\|_m \|\Lambda^{-1}\|_m, \end{aligned}$$

where $\|M\|_m = \min_{x: \|x\|_2=1} \|xM\|_2$. If $M \in \mathbb{R}^{a \times b}$, then $\|M\|_m$ is the a th largest singular value of M .

$$\begin{aligned} \|(\mu^y)^{-1}\|_m &= \|((Y^T Y)^{-1/2} U_{k_y})^{-1}\|_m = \|U_{k_y}^T (Y^T Y)^{1/2}\|_m \\ &\geq \|U_{k_y}^T\|_m \|(Y^T Y)^{1/2}\|_m \\ &= \|(Y^T Y)^{1/2}\|_m \end{aligned}$$

This is equal to the square root of the smallest block population in Y . Call this quantity $\sqrt{P_{\min}^y}$. The other important quantity, $\|\Lambda^{-1}\|_m$, is $1/\sigma_1$ where σ_1 is the largest singular value of \mathcal{L} . This is equal to 1. To see this, notice that $\mathcal{L}\mathcal{L}^T = \mathcal{P}^{-1/2} \mathcal{A} \mathcal{O}^{-1} \mathcal{A} \mathcal{P}^{-1/2}$ and the row sums of $\mathcal{A} \mathcal{O}^{-1} \mathcal{A}$ are contained down the diagonal of \mathcal{P} . So, $\mathcal{L}\mathcal{L}^T$ can be thought of as a classical symmetric graph Laplacian (from symmetric graphs). The largest eigenvalue of a symmetric graph Laplacian is 1 [von Luxburg, 2007]. This implies that $\|\Lambda^{-1}\|_m = 1$. Putting the pieces together,

$$(C.6) \quad \|\mathcal{V}_i - \mathcal{V}_j\|_2 \geq \sqrt{P_{\min}^y} \|B_{\cdot u}^{\mathcal{L}} - B_{\cdot v}^{\mathcal{L}}\|_2.$$

So, if the columns of $B^{\mathcal{L}}$ are unique and $u \neq v$, then $\|\mathcal{V}_i - \mathcal{V}_j\|_2 > 0$. The contrapositive implies that $\mathcal{V}_i = \mathcal{V}_j \Rightarrow z_i = z_j$. \square

Equivalence statements (C.4) and (C.5) imply that, under the Stochastic Co-Blockmodel with certain conditions, there are k_y unique rows in the right singular vectors, $Y\mu^y$, of \mathcal{L} and k_z unique rows in the left singular vectors, $Z\mu^z$ of \mathcal{L} . This has important consequences for DI-SIM. DI-SIM applied to \mathcal{L} will run k -means on the rows of $Y\mu^y$. Because there are only k_y unique points, each of these points will be a centroid of one of the resulting clusters. So, if $y_i\mu^y = y_j\mu^y$, then i and j will be assigned to the same cluster. With equivalence statement (C.4), this implies that DI-SIM applied to the matrix \mathcal{L} can perfectly identify the block memberships in Y . Similarly, DI-SIM applied to the matrix \mathcal{L} can perfectly identify the block memberships in Z . Obviously, \mathcal{L} is not observed. We need to estimate these memberships from the matrix L .

C.2. Comparing the population and observed clusters. Let $U \in \mathbb{R}^{n \times k_y}$ be a matrix whose orthonormal columns are the right singular vectors corresponding to the largest k_y singular values of L . DI-SIM applies k -means (with k_y clusters) to the rows of U , the points u_1, \dots, u_n . Each row is assigned to one cluster and each cluster has a centroid.

Definition 7. For $i = 1, \dots, n$, define $c_i^u \in \mathbb{R}^{k_y}$ to be the centroid corresponding to u_i .

Recall that $y_i\mu^y$ is the centroid corresponding to node i from the population analysis. If the observed centroid c_i^u is closer to the population centroid $y_i\mu^y$ than it is to any other population centroid $y_j\mu^y$ for $y_j \neq y_i$, then it appears that node i is correctly clustered. This definition is appealing because it removes one aspect of the cluster identifiability problem; instead of labeling the clusters $1, \dots, k_y$, the clusters are labeled by points in euclidean

space (the rows of μ^y). Unfortunately, the singular vectors impart one additional source of unidentifiability. They are only identifiable up to a rotation. The orthonormal rotation from Theorem B.2, $\mathcal{R} \in \mathbb{R}^{k_y \times k_y}$, can address this technical nuisance. Consider node i to be correctly clustered if, c_i^u is closer to $y_i \mu^y \mathcal{R}$ than it is to any other (rotated) population centroid $y_j \mu^y \mathcal{R}$ for $y_j \neq y_i$. The slight complication with \mathcal{R} stems from the fact that the vectors c_1^u, \dots, c_n^u are constructed from the singular vectors in U and Theorem B.2 in the Appendix shows that, under certain conditions, the singular vectors of L converge to an orthonormal rotation of the singular vectors of \mathcal{L} . In other words, the singular vectors U converge to the *rotated* population eigenvectors: $Y \mu^y \mathcal{R}$.

Define P_{max}^y to be the population of the largest block in Y .

$$(C.7) \quad P_{max}^y = \max_{j=1, \dots, k_y} (Y^T Y)_{jj}$$

The following lemma motivates the definition of misclustered by providing a sufficient condition for a node to be correctly clustered.

Lemma C.2. *For any orthonormal matrix $\mathcal{R} \in \mathbb{R}^{k_y \times k_y}$,*

$$(C.8) \quad \|c_i^u \mathcal{R} - y_i \mu^y\|_2 < 1/\sqrt{2P_{max}^y} \implies$$

$$(C.9) \quad \|c_i^u \mathcal{R} - y_i \mu^y\|_2 < \|c_i^u \mathcal{R} - y_j \mu^y\|_2 \text{ for any } y_j \neq y_i.$$

Proof. Statement (C.10) is the essential ingredient to prove Lemma C.2.

$$(C.10) \quad y_i \neq y_j, \text{ then } \|y_i \mu^y - y_j \mu^y\|_2 \geq \sqrt{2/P_{max}^y}$$

The proof of statement (C.10) requires the following definition,

$$\|\mu^y\|_m^2 = \min_{x: \|x\|_2=1} \|x \mu^y\|_2^2.$$

Notice that

$$\|\mu^y\|_m^2 = \min_{x: \|x\|_2=1} x \mu^y (\mu^y)^T x^T = \min_{x: \|x\|_2=1} x (Y^T Y)^{-1} x^T = 1/P_{max}^y.$$

So,

$$\|y_i \mu^y - y_j \mu^y\|_2 = \|(y_i - y_j) \mu^y\|_2 \geq \sqrt{2} \|\mu^y\|_m = \sqrt{2/P_{max}^y},$$

Proving statement (C.10). The proof of Lemma C.2 follows,

$$\|c_i^u \mathcal{R} - y_j \mu^y\|_2 \geq \|y_i \mu^y - y_j \mu^y\|_2 - \|c_i^u \mathcal{R} - y_i \mu^y\|_2 \geq \sqrt{\frac{2}{P_{max}^y}} - \frac{1}{2} \sqrt{\frac{2}{P_{max}^y}} = \frac{1}{\sqrt{2P_{max}^y}}$$

□

Line (C.9) is the previously motivated definition of correctly clustered. Thus, Lemma C.2 shows that inequality (C.8) is a sufficient condition for node i to be correctly clustered in the Y block.

Definition 8. For the orthonormal rotation \mathcal{R} from Proposition B.1, define the set of y-misclustered nodes as the nodes that do not satisfy (C.8),

$$(C.11) \quad \mathcal{M}_y = \left\{ i : \|c_i^u \mathcal{R} - y_i \mu^y\|_2 \geq 1/\sqrt{2P_{max}^y} \right\}.$$

Thus far, this subsection has motivated a definition of y-misclustered. Because $k_y \leq k_z$, it is slightly more challenging to give a definition for z-misclustered.

Let $V \in \mathbb{R}^{n \times k_y}$ be a matrix whose orthonormal columns are the left singular vectors corresponding to the largest k_y singular values of L .

Definition 9. For $i = 1, \dots, n$, define $c_i^v \in \mathbb{R}^{k_y}$ to be the centroid corresponding to V_i .

To give the analogue to Lemma C.2, define

$$P_{min}^y = \min_{j=1, \dots, k_y} (Y^T Y)_{jj},$$

define $B_{\cdot j}^{\mathcal{L}}$ as the j th column of $B^{\mathcal{L}}$, and define

$$(C.12) \quad \gamma_z = \sqrt{P_{min}^y} \min_{i \neq j} \|B_{\cdot i}^{\mathcal{L}} - B_{\cdot j}^{\mathcal{L}}\|_2.$$

Lemma C.3. For any orthonormal matrix $\mathcal{R} \in \mathbb{R}^{k_y \times k_y}$,

$$(C.13) \quad \|c_i^v \mathcal{R} - z_i \mu^z\|_2 < \gamma_z/2 \implies$$

$$(C.14) \quad \|c_i^v \mathcal{R} - z_i \mu^z\|_2 < \|c_i^v \mathcal{R} - z_j \mu^z\|_2 \text{ for any } z_j \neq z_i.$$

Proof. If $\mathcal{V}_i \neq \mathcal{V}_j$, then using equation C.6,

$$\|z_i \mu^z - z_j \mu^z\|_2 = \|\mathcal{V}_i - \mathcal{V}_j\|_2 \geq \sqrt{P_{min}^y} \|B_{\cdot u}^{\mathcal{L}} - B_{\cdot v}^{\mathcal{L}}\|_2 \geq \gamma_z$$

and the proof of Lemma C.3 follows,

$$\|c_i^v \mathcal{R} - z_j \mu^z\|_2 \geq \|z_i \mu^z - z_j \mu^z\|_2 - \|c_i^v \mathcal{R} - z_i \mu^z\|_2 > \gamma_z - \gamma_z/2 = \gamma_z/2 > \|c_i^v \mathcal{R} - z_i \mu^z\|_2.$$

□

Definition 10. For the orthonormal rotation \mathcal{R} from Proposition B.1, define the set of z-misclustered nodes as the nodes that do not satisfy (C.13),

$$(C.15) \quad \mathcal{M}_z = \{i : \|c_i^v \mathcal{R} - z_i \mu^z\|_2 \geq \gamma_z/2\}.$$

Proof. This is a proof of Theorem 4.1. It follows from Lemma C.1 that $\mathcal{L}^{(n)}$ has k_y nonzero singular values. In order to invoke Theorem B.2, define $S_n = [\sigma_{k_y}^2/2, \infty)$. Similar to the argument in the proof of Theorem B.2, by Weyl's inequality, Theorem B.1, and the assumptions of Theorem 4.1, the top k_y singular values of $\mathcal{L}^{(n)}$ will eventually be contained in S_n . This definition of S_n implies that $\delta_n = \delta'_n = \sigma_{k_y}^2/2$.

Define a partition matrix to be a matrix of all zeros except for a single 1 in each row. Define

$$\mathcal{R}(n, k) = \left\{ \tilde{Y}T : \text{ where } \tilde{Y} \in \mathbb{R}^{n \times k} \text{ is a partition matrix and } T \in \mathbb{R}^{k \times k} \right\}.$$

Notice that

$$\min_{M \in \mathcal{R}(n, k)} \|U - M\|_F^2 = \min_{\{m_1, \dots, m_k\} \subset \mathbb{R}^k} \sum_i \min_g \|u_i - m_g\|_2^2$$

where the right hand side of the equation is the k -means objective function applied to the rows of V . Define

$$C_U = \operatorname{argmin}_{M \in \mathcal{R}(n, k)} \|U - M\|_F^2.$$

Notice that the i th row of C_U is c_i^u . Further, $C_U \mathcal{R} = \operatorname{argmin}_{M \in \mathcal{R}(n, k)} \|U \mathcal{R} - M\|_F^2$. Because $Y \mu^y \in \mathcal{R}(n, k)$,

$$\|C_U \mathcal{R} - Y \mu^y\|_F \leq \|C_U \mathcal{R} - U \mathcal{R}\|_F + \|U \mathcal{R} - Y \mu^y\|_F \leq 2\|U \mathcal{R} - Y \mu^y\|_F.$$

Therefore,

$$\begin{aligned} |\mathcal{M}_y| \leq \sum_{i \in \mathcal{M}_y} 1 &\leq 2P_{max}^y \sum_{i \in \mathcal{M}_y} \|c_i^u \mathcal{R} - y_i \mu^y\|_2^2 \\ &\leq 2P_{max}^y \|C_U \mathcal{R} - Y \mu^y\|_F^2 \\ &\leq 4P_{max}^y \|U \mathcal{R} - Y \mu^y\|_F^2 \\ &= o\left(\frac{P_{max}^y (\log n)^2}{n \sigma_{k_y}^4 \tau_n^4}\right) \end{aligned}$$

The last line follows from Theorem B.2.

Similar statements hold for Z . Define

$$C_V = \operatorname{argmin}_{M \in \mathcal{R}(n, k)} \|V - M\|_F^2.$$

Thus, similar to before,

$$\|C_V \mathcal{R} - Z \mu\|_F \leq 2\|V \mathcal{R} - Z \mu^z\|_F.$$

Therefore,

$$\begin{aligned} |\mathcal{M}_z| \leq \sum_{i \in \mathcal{M}_z} 1 &\leq 4\gamma_z^{-2} \sum_{i \in \mathcal{M}_z} \|c_i^v \mathcal{R} - z_i \mu^z\|_2^2 \\ &\leq 4\gamma_z^{-2} \|C_V \mathcal{R} - Z \mu^z\|_F^2 \\ &\leq 8\gamma_z^{-2} \|V \mathcal{R} - Z \mu^z\|_F^2 \\ &= o\left(\frac{(\log n)^2}{n \sigma_{k_y}^4 \tau_n^4 \gamma_z^2}\right) \end{aligned}$$

□

Next is a proof of Proposition 4.1.

Proof. To simplify

$$\gamma_z = \sqrt{P_{\min}^y} \min_{i \neq j} \|B_{\cdot i}^{\mathcal{L}} - B_{\cdot j}^{\mathcal{L}}\|_2,$$

first simplify $B^{\mathcal{L}} = D_y^{-1/2} B D_z^{-1/2}$. The (i, i) th element of $D_z^{-1/2}$ is

$$\frac{1}{\sqrt{\mathbf{1}_n^T Y B_{\cdot i}}} = \frac{1}{\sqrt{n\tau_n}}.$$

Similarly for $D_y^{-1/2}$. So,

$$B^{\mathcal{L}} = \frac{1}{n\tau_n} B.$$

Thus,

$$\gamma_z = \frac{\sqrt{P_{\min}^y}}{n\tau_n} \min_{i \neq j} \|B_{\cdot i} - B_{\cdot j}\|_2.$$

□

The following is a proof of Corollary 4.1

Proof. From the proof of Lemma C.1, $\mathcal{L}^T \mathcal{L}$ has the same eigenvalues as $(Z^T Z)^{1/2} B_{LL} (Z^T Z)^{1/2}$ and these values are the squared singular values of \mathcal{L} . Under the four parameter Stochastic Co-Blockmodel,

$$(Z^T Z)^{1/2} B_{LL} (Z^T Z)^{1/2} = \frac{1}{(kr + p)^2} (p^2 I_k + (2p + r)r \mathbf{1}_k \mathbf{1}_k^T).$$

The constant vector is an eigenvector of this matrix. It has eigenvalue

$$\sigma_1^2 = \frac{p^2 + 2kpr + kr^2}{(kr + p)^2}.$$

Any vector orthogonal to a constant vector is also an eigenvector. They all have eigenvalue

$$\sigma_{k_y}^2 = \frac{p^2}{(kr + p)^2} = \frac{1}{(k(r/p) + 1)^2}.$$

Notice that $\tau > r$. If $k = O(n^{1/4}/\log n)$, then the conditions of Theorem 4.1 are satisfied. Substituting the values of $P = s$ and σ_{k_y} above,

$$|\mathcal{M}_y| = o\left(\frac{s(\log n)^2 k^4}{n}\right) = o(k^3(\log n)^2).$$

Because $k_y = k_z$, \mathcal{M}_z can be defined as

$$\mathcal{M}_z = \left\{ i : \|c_i^v \mathcal{R} - z_i \mu^z\|_2 \geq 1/\sqrt{2P_{\max}^z} \right\},$$

where $P_{max}^z = s$. This definition is analogous the definition of \mathcal{M}_y (Equation C.11). The result for \mathcal{M}_y above also holds for \mathcal{M}_z . \square

REFERENCES

- A. Banerjee, I. Dhillon, J. Ghosh, S. Merugu, and D.S. Modha. A generalized maximum entropy approach to bregrman co-clustering and matrix approximation. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 509–514. ACM, 2004.
- M. Belkin. *Problems of learning on manifolds*. PhD thesis, The University of Chicago, 2003.
- M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008.
- R. Bhatia. *Perturbation Bounds for Matrix Eigenvalues (Classics in Applied Mathematics)*. Society for Industrial and Applied Mathematics, 2007. ISBN 0898716314.
- P.J. Bickel and A. Chen. A nonparametric view of network models and Newman–Girvan and other modularities. *Proceedings of the National Academy of Sciences*, 106(50):21068, 2009.
- G. Bisson and F. Hussain. Chi-sim: A new similarity measure for the co-clustering task. In *Machine Learning and Applications, 2008. ICMLA’08. Seventh International Conference on*, pages 211–217. IEEE, 2008.
- O. Bousquet, O. Chapelle, and M. Hein. Measure based regularization. *Advances in Neural Information Processing Systems*, 16, 2004.
- D.S. Choi, P.J. Wolfe, and E.M. Airolidi. Stochastic blockmodels with growing number of classes. *Biometrika (in press)*, 2012.
- F.R.K. Chung. *Spectral graph theory*. Amer Mathematical Society, 1997.
- R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- I.S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 269–274. ACM, 2001.
- W.E. Donath and A.J. Hoffman. Lower bounds for the partitioning of graphs. *IBM Journal of Research and Development*, 17(5):420–425, 1973.
- M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- D. Freitag. Trained named entity recognition using distributional clusters. In *Proceedings of EMNLP*, volume 4, pages 262–269, 2004.
- E. Giné and V. Koltchinskii. Empirical graph Laplacian approximation of Laplace–Beltrami operators: large sample results. *Lecture Notes–Monograph Series*, pages 238–259, 2006.

- Y. Goldberg, A. Zakai, and D. Kushnir. Manifold learning: The price of normalization. *The Journal of Machine Learning Research*, 9:1909–1939, 2008.
- L. Guttman. Metricizing rank-ordered or unordered data for a linear factor analysis. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 21(3/4):257–268, 1959. ISSN 0036-4452.
- L. Hagen and A.B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Trans. Computer-Aided Design*, 11(9):1074–1085, 1992.
- J.A. Hartigan. Direct clustering of a data matrix. *Journal of the American Statistical Association*, pages 123–129, 1972.
- M. Hein. Uniform convergence of adaptive graph-based regularization. *Lecture Notes in Computer Science*, 4005:50, 2006.
- M. Hein, J.Y. Audibert, and U. von Luxburg. Graph Laplacians and their convergence on random neighborhood graphs. *Arxiv preprint math.ST/0608522*, 2006.
- B. Hendrickson and R. Leland. An improved spectral graph partitioning algorithm for mapping parallel computations. *SIAM Journal on Scientific Computing*, 16(2):452–469, 1995.
- M.O. Hill. *DECORANA: a FORTRAN program for detrended correspondence analysis and reciprocal averaging*. Cornell University, 1979.
- HO Hirschfeld. A connection between correlation and contingency. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 31, pages 520–524. Cambridge Univ Press, 1935.
- P.D. Hoff. Multiplicative latent factor models for description and prediction of social networks. *Computational & Mathematical Organization Theory*, 15(4):261–272, 2009. ISSN 1381-298X.
- P.D. Hoff, A.E. Raftery, and M.S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002. ISSN 0162-1459.
- P. Holland, K.B. Laskey, and S. Leinhardt. Stochastic blockmodels: Some first steps. *Social Networks*, 5:109–137, 1983.
- S. Holmes. Multivariate analysis: The french way. *Festschrift for David Freedman*, 2006.
- R.A. Horn and C.R. Johnson. *Matrix analysis*. Cambridge university press, 2005. ISBN 0521386322.
- J.M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- Vladimir Koltchinskii and Evarist Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, 6(1):113–167, 2000. ISSN 13507265. URL <http://www.jstor.org/stable/3318636>.
- J. Leskovec, K.J. Lang, A. Dasgupta, and M.W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proceeding of the 17th international conference on World Wide Web*, pages 695–704. ACM, 2008.

- S.C. Madeira and A.L. Oliveira. Biclustering algorithms for biological data analysis: a survey. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 1(1): 24–45, 2004.
- S.C. Madeira, M.C. Teixeira, I. Sa-Correia, and A.L. Oliveira. Identification of regulatory modules in time series gene expression data using a linear time biclustering algorithm. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 7(1):153–165, 2010.
- L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. 1999.
- A. Pothen, H.D. Simon, and K.P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM Journal on Matrix Analysis and Applications*, 11:430, 1990.
- K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- R. Rohwer and D. Freitag. Towards full automation of lexicon construction. In *Proceedings of the HLT-NAACL Workshop on Computational Lexical Semantics*, pages 9–16. Association for Computational Linguistics, 2004.
- V. Satuluri and S. Parthasarathy. Symmetrizations for clustering directed graphs. In *Proceedings of the 14th International Conference on Extending Database Technology*, pages 343–354. ACM, 2011.
- J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- T. Shi, M. Belkin, and B. Yu. Data Spectroscopy: Eigenspace of Convolution Operators and Clustering. *The Annals of Statistics*, 37(6B):3960–3984, 2009.
- D.A. Spielman and S.H. Teng. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421(2-3):284–305, 2007.
- H. Steinhaus. Sur la division des corp materiels en parties. *Bull. Acad. Polon. Sci*, 1: 801–804, 1956.
- D.L. Sussman, M. Tang, D.E. Fishkind, and C.E. Priebe. A consistent dot product embedding for stochastic blockmodel graphs. *Arxiv preprint arXiv:1108.2228*, 2011.
- A. Tanay, R. Sharan, M. Kupiec, and R. Shamir. Revealing modularity and organization in the yeast molecular network by integrated analysis of highly heterogeneous genomewide data. *Proceedings of the National Academy of Sciences of the United States of America*, 101(9):2981, 2004.
- A. Tanay, R. Sharan, and R. Shamir. Biclustering algorithms: A survey. *Handbook of computational molecular biology*, 9:26–1, 2005.
- C.J.F. Ter Braak. Canonical correspondence analysis: a new eigenvector technique for multivariate direct gradient analysis. *Ecology*, 67(5):1167–1179, 1986. ISSN 0012-9658.
- D Ting, L Huang, and M Jordan. An analysis of the Convergence of the Graph Laplacians. *Arxiv preprint math.ML/1101.5435v1*, 2011.

- R. Van Driessche and D. Roose. An improved spectral bisection algorithm and its application to dynamic load balancing. *Parallel Computing*, 21(1):29–48, 1995.
- U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.
- U. von Luxburg, M. Belkin, and O. Bousquet. Consistency of spectral clustering. *Annals of Statistics*, 36(2):555, 2008.
- Y.J. Wang and G.Y. Wong. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987. ISSN 0162-1459.
- Y. Zhao, E. Levina, and J. Zhu. On consistency of community detection in networks. *Arxiv preprint arXiv:1110.3854*, 2011.